

Ulrike Spierling
Nicolas Szilas (Eds.)

LNCS 5334

Interactive Storytelling

First Joint International Conference
on Interactive Digital Storytelling, ICIDS 2008
Erfurt, Germany, November 2008, Proceedings

 Springer

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Ulrike Spierling Nicolas Szilas (Eds.)

Interactive Storytelling

First Joint International Conference
on Interactive Digital Storytelling, ICIDS 2008
Erfurt, Germany, November 26-29, 2008
Proceedings

Volume Editors

Ulrike Spierling
FH Erfurt, University of Applied Sciences
Altonaer Straße 25, 99085 Erfurt, Germany
E-mail: spierling@fh-erfurt.de

Nicolas Szilas
Université de Genève
Faculté de Psychologie et des Sciences de l'Éducation
TECFA
1211 Genève 4, Switzerland
E-mail: Nicolas.Szilas@unige.ch

Library of Congress Control Number: 2008939012

CR Subject Classification (1998): H.4, I.3, I.2, C.3, H.5, I.4, I.7.2, J.5

LNCS Sublibrary: SL 3 – Information Systems and Application, incl. Internet/Web and HCI

ISSN 0302-9743
ISBN-10 3-540-89424-1 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-89424-7 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2008
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12567485 06/3180 5 4 3 2 1 0

Preface

This volume contains scientific papers and case studies presented at Interactive Storytelling '08: The First Joint International Conference on Interactive Digital Storytelling (ICIDS), held November 26–29, 2008, in Erfurt, Germany.

Interactive Digital Storytelling (IDS) is a cross-disciplinary topic, which explores new uses of interactive technologies for creating and experiencing narratives. IDS is also a huge step forward in games and learning. This can be seen through its ability to enrich virtual characters with intelligent behavior, to allow collaboration of humans and machines in the creative process, and to combine narrative knowledge and user activity in interactive artifacts.

IDS involves concepts from many aspects of Computer Science, above all from Artificial Intelligence, with topics such as narrative intelligence, automatic dialogue and drama management, and smart graphics. In order to process stories in real time, traditional storytelling needs to be formalized into computable models by drawing from narratological studies. As it is currently hardly accessible for creators and end-users, there is a need for new authoring concepts and tools supporting the creation of such dynamic stories, allowing for rich and meaningful interaction with the content.

On the past 7 years, two European conference series have been addressing these problems, serving as annual meeting points for scientists, researchers and developers: ICVS (International Conference on Virtual Storytelling) and TIDSE (Technologies for Interactive Digital Storytelling and Entertainment) took place bi-annually and alternated between France and Germany. The year 2008 was an opportunity for consolidation: The organizers of both lines decided to integrate them and constitute a joint international series of annual conferences, addressing all the topical research issues in conjunction with the themes of the previous conferences.

This first conference, Interactive Storytelling '08, received great attention in the research community. In all, 42 Program Committee members selected 19 full papers for publication out of 62 submissions – with an acceptance rate of 0.3. Further, this volume has been completed by five short papers, five poster presentations and eight demonstrations of interactive storytelling applications and case studies. Many contributions present architectures for automatic story generation and drama management, for virtual characters and new authoring tools. Some relatively new topics include the evaluation of dramatic immersion and music.

The program also included three invited talks positing future perspectives of the field. Additional workshops on pen and paper role playing and improvisational theater were scheduled, providing opportunities to learn from non-digital forms of interactive storytelling. Two panels discussing industrial applications and issues of creation also addressed the still recognizable gap between available technologies and their accessibility for creation.

Has the field reached the ‘*âge de raison*’ – a French expression (‘age of rationality’) denoting the fact that children reach a certain maturity at the age of seven? Do we foresee signs of maturity in IDS? Now, with a consolidated annual event, we face new

challenges that will lead us into adolescence: The novel technologies will only succeed if the conflicts between the disciplines of technical sciences and of humanities can be bridged. Our field of IDS benefits from new ideas arising from the combination of computational rationality with the ‘unreasonableness’ of the creative spirit, which is at the heart of storytelling. The new conference series – by continuing the work of its predecessors – shall provide opportunities for growing up. We hope that participants in this first venue found inspiration and guidelines for future research.

We want to thank the group of previous ICVS / TIDSE organizers for the constitution of this new conference line. Foremost, we have to thank the amazing members of our Program Committee for their sound and high-quality reviews, which contributed significantly to the success of this conference. Finally, we want to thank all sponsors and collaborators for their support.

November 2008

Ulrike Spierling
Nicolas Szilas

Organization

Conference Chairs

Ulrike Spierling FH Erfurt, University of Applied Sciences, Erfurt,
Germany
Nicolas Szilas TECFA, University of Geneva, Switzerland

Program Committee

Elisabeth André University of Augsburg, Germany
Ruth Aylett Heriot-Watt University, UK
Jay David Bolter Center for New Media Research and Education,
Atlanta, USA
Brunhild Bushoff sagasnet, Munich, Germany
Marc Cavazza University of Teesside, Middlesbrough, UK
Ronan Champagnat L3i, IUT de La Rochelle, France
Adrian David Cheok National University of Singapore, Singapore
Chris Crawford Storytron.com, USA
Stéphane Donikian IRISA / INRIA, Rennes, France
Stefan Göbel ZGDV e.V., Darmstadt, Germany
Andrew Gordon USC Institute for Creative Technology, Marina del
Rey, USA
Paul Grimm FH Erfurt University of Applied Sciences, Erfurt,
Germany
Holger Großmann Fraunhofer IDMT, Ilmenau, Germany
Knut Hartmann Flensburg University of Applied Sciences, Flensburg,
Germany
Ido Iurgel CCG, Centro de Computação Gráfica, Guimarães,
Portugal
Klaus P. Jantke Fraunhofer IDMT, Ilmenau, Germany
Craig Lindley Blekinge Technical College, Sweden
Brian Magerko Georgia Institute of Technology, Atlanta, USA
Stacy Marsella USC Information Sciences Institute, Marina Del Rey,
USA
Maic Masuch Trier University of Applied Sciences, Trier, Germany
Wolfgang Müller University of Education, Weingarten, Germany
Ryohei Nakatsu National University of Singapore, Singapore
Stephane Natkin Cedric, France
Michael Nitsche Georgia Institute of Technology, Atlanta, USA
Ana Paiva INESC-ID, Portugal
Zhigeng Pan Zhejiang University, Hangzhou, China

Paolo Petta	Austrian Research Institute for Artificial Intelligence, Vienna, Austria
Yusuf Pisan	University of Technology, Sydney, Australia
Jean-Hugues Réty	University of Paris 8, Saint-Denis, France
Mark Riedl	Georgia Institute of Technology, Atlanta, USA
Leonie Schäfer	European Commission, Brussels, Belgium
Oliver Schneider	IGDV, University of Applied Sciences, Darmstadt, Germany
Magy Seif El-Nasr	Simon Fraser University, Surrey, Canada
Thomas Sporer	Fraunhofer IDMT, Ilmenau, Germany
Andrew Stern	Procedural Arts, LLC, Portland, USA
Wen-Poh Su	Griffith University, Brisbane, Australia
Kaoru Sumi	National Institute of Information & Communications Technology, Kyoto, Japan
Richard Wages	NOMADS Lab, Cologne, Germany
R. Michael Young	North Carolina State University, Raleigh, USA
Nelson Zagalo	University of Minho, Braga, Portugal

Further Reviewers: Daniel Gärtner, Isuru Godage, Vinod Krishnan, Hanna Lukashevich, Nimesha Ranasinghe, Stefan Rank, David Roberts, Sebastian Weiß, Kening Zhu.

Constituting Committee of the Joint Interactive Storytelling Conference Series: Olivier Balet, Marc Cavazza, Ronan Champagnat, Stéphane Donikian, Stefan Göbel, Ido Jurgel, Ulrike Spierling, Nicolas Szilas.

Local Organization: Interactive Storytelling '08 was organized by the Fachhochschule Erfurt, University of Applied Sciences, in Erfurt, Germany, in collaboration with the following institutions: 'Das KinderMedienZentrum' in Erfurt, 'Thüringer Staatskanzlei' (Thuringia State Chancellery), and the Erfurt project group 'Children's Media' of the Fraunhofer IDMT in Ilmenau. Special thanks to Paul Grimm, FH Erfurt, and Klaus Peter Jantke, Holger Großmann and Thomas Sporer, Fraunhofer IDMT, for their collaboration. Further sponsors and supporters are listed on the conference website.

Conference Website: <http://www.fh-erfurt.de/ai/icids08>

This proceedings volume has been supported by the Thuringia foundation for technology, innovation and research (Stiftung für Technologie, Innovation und Forschung Thüringen).



Stiftung für Technologie,
Innovation und Forschung
Thüringen

Table of Contents

Future Perspectives on Interactive Digital Storytelling (Keynotes)

Embracing the Combinatorial Explosion: A Brief Prescription for Interactive Story R&D	1
<i>Andrew Stern</i>	
Interactive Narrative, Plot Types, and Interpersonal Relations	6
<i>Marie-Laure Ryan</i>	
The IRIS Network of Excellence: Integrating Research in Interactive Storytelling	14
<i>Marc Cavazza, Stéphane Donikian, Marc Christie, Ulrike Spierling, Nicolas Szilas, Peter Vorderer, Tilo Hartmann, Christoph Klimmt, Elisabeth André, Ronan Champagnat, Paolo Petta, and Patrick Olivier</i>	

Interactive Storytelling Applications

Mobile Urban Drama – Setting the Stage with Location Based Technologies	20
<i>Frank Allan Hansen, Karen Johanne Kortbek, and Kaj Grønbaek</i>	
Say Anything: A Massively Collaborative Open Domain Story Writing Companion	32
<i>Reid Swanson and Andrew S. Gordon</i>	
Locating Drama: A Demonstration of Location-Aware Audio Drama . . .	41
<i>Nye Parry, Helen Bendon, Stephen Boyd Davis, and Magnus Moar</i>	
Lies and Seductions	44
<i>Petri Lankoski and Tommi Horttana</i>	
Animation-Based Interactive Storytelling System	48
<i>Kaoru Sumi</i>	
Dear Esther: An Interactive Ghost Story Built Using the Source Engine	51
<i>Dan Pinchbeck</i>	
Walking the Edit – A Research Project of the Master Cinema Network in Switzerland	55
<i>Ulrich Fischer</i>	

Virtual Characters and Agents

3D Immersion in Virtual Agents Education	59
<i>Cyril Brom, Jakub Gemrot, Ondřej Burkert, Rudolf Kadlec, and Michal Bída</i>	
Exploring Non-verbal Behavior Models for Believable Characters	71
<i>Magy Seif El-Nasr and Huaxin Wei</i>	
Revisiting Character-Based Affective Storytelling under a Narrative BDI Framework	83
<i>Federico Peinado, Marc Cavazza, and David Pizzi</i>	
<i>VirtualActor</i> : Endowing Virtual Characters with a Repertoire for Acting	89
<i>Ido A. Iurgel</i>	
Steps towards a Generic Interface between Interactive Storytelling Applications and Character Animation Engines	92
<i>Sebastian A. Weiß, Florian Berger, Alexander Marbach, and Wolfgang Müller</i>	

User Experience and Dramatic Immersion

Looking at the Interactive Narrative Experience through the Eyes of the Participants	96
<i>David Milam, Magy Seif El-Nasr, and Ron Wakkary</i>	
Play and Narration as Patterns of Meaning Construction: Theoretical Foundation and Empirical Evaluation of the User Experience of Interactive Films	108
<i>Regina Friess</i>	
Trying to Get Trapped in the Past – Exploring the Illusion of Presence in Virtual Drama	114
<i>Georg Struck, Ralf Böse, and Ulrike Spierling</i>	
The Functions of Music in Interactive Media	126
<i>Axel Berndt and Knut Hartmann</i>	
Adaptive Musical Expression from Automatic Realtime Orchestration and Performance	132
<i>Axel Berndt and Holger Theisel</i>	

Architectures for Story Generation

Narrative Generation for Suspense: Modeling and Evaluation	144
<i>Yun-Gyung Cheong and R. Michael Young</i>	

A Use of Flashback and Foreshadowing for Surprise Arousal in Narrative Using a Plan-Based Approach	156
<i>Byung-Chull Bae and R. Michael Young</i>	
Story Planning with Vignettes: Toward Overcoming the Content Production Bottleneck	168
<i>Mark O. Riedl and Neha Sugandh</i>	
Schemas in Directed Emergent Drama	180
<i>Maria Arinbjarnar and Daniel Kudenko</i>	
Developing a Drama Management Architecture for Interactive Fiction Games	186
<i>Santiago Ontañón, Abhishek Jain, Manish Mehta, and Ashwin Ram</i>	
Planning and Interaction Levels for TV Storytelling	198
<i>Angelo E.M. Ciarlini, Marcelo M. Camanho, Thiago R. Dória, Antonio L. Furtado, Cesar T. Pozzer, and Bruno Feijó</i>	
Exploiting Structure and Conventions of Movie Scripts for Information Retrieval and Text Mining	210
<i>Arnav Jhala</i>	
Generation of Dilemma-Based Narratives: Method and Turing Test Evaluation	214
<i>Heather Barber and Daniel Kudenko</i>	
Models for Drama Management and Interacting with Stories	
Emergent Stories Facilitated: An Architecture to Generate Stories Using Intelligent Synthetic Characters	218
<i>Rui Figueiredo, Antonio Brisson, Ruth Aylett, and Ana Paiva</i>	
Making Stories Player-Specific: Delayed Authoring in Interactive Storytelling	230
<i>David Thue, Vadim Bulitko, and Marcia Spetch</i>	
Verbal Communication of Story Facilitators in Multi-player Role-Playing Games	242
<i>Anders Tychsen, Thea Brolund, and Michael Hitchens</i>	
Improvisation and Performance as Models for Interacting with Stories	250
<i>Theresa Jean Tanenbaum and Karen Tanenbaum</i>	
Let's Pretend I Had a Sword: Late Commitment in Emergent Narrative	264
<i>Ivo Swartjes, Edze Kruizinga, and Mariët Theune</i>	

On the Use of Computational Models of Influence for Managing
Interactive Virtual Experiences 268
*David L. Roberts, Charles Isbell, Mark Riedl, Ian Bogost, and
Merrick L. Furst*

Authoring and Creation of Interactive Narratives

Purposeful Authoring for Emergent Narrative 273
Sandy Louchart, Ivo Swartjes, Michael Kriegel, and Ruth Aylett

From Debugging to Authoring: Adapting Productivity Tools to
Narrative Content Description 285
David Pizzi and Marc Cavazza

PRISM: A Framework for Authoring Interactive Narratives 297
*Yun-Gyung Cheong, Yeo-Jin Kim, Wook-Hee Min,
Eok-Soo Shim, and Jin-Young Kim*

Tales for the Many: Process and Authorial Control in Multi-player
Role-Playing Games 309
Anders Tychsen

An Intelligent Plot-Centric Interface for Mastering Computer
Role-Playing Games 321
*Carlos León, Federico Peinado, Álvaro Navarro, and
Héctor Cortiguera*

StoryTec: A Digital Storytelling Platform for the Authoring and
Experiencing of Interactive and Non-linear Stories 325
*Stefan Göbel, Luca Salvatore, Robert Arthur Konrad, and
Florian Mehm*

Workshop: Impro Theatre 329
Dagmar Dörger and Martin Geisler

Workshop: Pen-and-Paper Role-Playing 330
Florian Berger and Alexander Marbach

Workshop and Panel: The Authoring Process in Interactive
Storytelling 331
Ulrike Spierling and Ido Iurgel

Author Index 333

Embracing the Combinatorial Explosion: A Brief Prescription for Interactive Story R&D

Andrew Stern

Procedural Arts
107 SE Washington St, Suite 450
Portland, Oregon, USA
andrew@proceduralarts.com

Abstract. In this editorial, the quintessential requirements for successful interactive story systems are examined and discussed. As part of the discussion, a brief prescription for interactive story research and development is suggested.

Keywords: Interactive story, art, agency, generation, interface, drama management.

1 Unreached Potential

As a new form of experiencing narrative, digital interactive story has the potential to become a premier artform of the 21st century. However, this potential is still far from being realized. While the challenge that researchers and developers face in this endeavor is primarily a technological one, it is also deeply intertwined with design.

We find it useful to compare the development of interactive story to other endeavors that have both similar and different characteristics, to help us understand the special nature of the challenge and what research and development it may take to accomplish it.

1.1 Comparison to Mainstream AI Technology

Universal speech recognition, including language understanding, is one of those potential mainstream AI technologies that could revolutionize society's interactions with computers, akin to the kind of societal impact that some researchers and developers believe interactive story could have. As it turns out, both endeavors have been far more technically daunting to accomplish than originally hoped, and require several support technologies that are still underdeveloped. For example, each can greatly benefit from, and in fact may require, broad common-sense knowledge and reasoning about the world; this piece of the puzzle is a huge technical challenge in its own right.

Unlike speech recognition though, interactive story is more than just a technology: it is first and foremost a creative endeavor, an artform. Successful interactive stories will require artists to harness the underlying technology to produce satisfying, enriching and meaningful experiences for players. The challenge for creating successful interactive stories can be seen as even greater, or at least broader and more interdisciplinary, than that of mainstream AI technologies alone, such as speech recognition.

To put it further into perspective, R&D into speech recognition has been ongoing for longer and has received more funding than R&D into interactive story. And while speech recognition R&D has made some good progress, with some applications in use today, it has only been partially accomplished. The interactive story research community's expectations about the challenges it faces should be calibrated accordingly.

1.2 Comparison to Cinema

Interactive story as a narrative form can be said to be in an infant stage of development, akin to the first cinematic films of the early 20th century. In some respects, the comparison seems apt. The birth of cinema saw the application of impressive new technologies, particularly the ability to photograph and record motion and replay it later for large audiences; this was followed on by development of techniques of editing, montage, cinematography, etc., and more generally, innovations in narrative structure and story representation. We can see parallels in the development of the grammar of digital interactive stories of today.

However, the technological challenges of developing cinema, while significant at the time, were arguably an order of magnitude simpler than those required for digital interactive story today. In fact, the subset of technologies directly akin to cinema required for interactive story already exist, specifically, the ability to animate and visually present human behavior and its environment. But interactive story is fundamentally about interaction, and the system intelligence to support that. Dynamic simulation of human behavior is far more complex to engineer than the static representation of it. Therefore, we cannot expect the development of interactive story to progress at the same rate as cinema did a century ago; the technological hurdles are very different.

1.3 Comparison to Videogames

Perhaps the most appropriate comparison of interactive story is to videogames. Today's games have had a strong impact on society's entertainment landscape, even if they are still mostly considered an action-oriented, often juvenile entertainment form. Building games requires a combination and balance of technology and design in similar proportion to what is required for interactive story – a closer analogy than speech recognition or cinema.

There are stories in today's games, of course, but their overarching plots are invariably linear, and therefore cannot be considered fully interactive (contrary to some game developer's hyperbolic claims). Of the games that offer more open-ended, non-linear gameplay, they offer fragments of coherent emergent story at best.

That said, there are characteristics of games that have made them very successful in the marketplace to date, that will be fundamentally critical to the potential of interactive stories as well. Putting aside superficial features such as high fidelity graphics and animation, the most successful games offer players true *agency* – the ability for players to have persistent, meaningful effects on the events of the experience.

Today's games have only achieved agency in the domain of action- or puzzle-oriented gameplay, but we believe that has been enough to demonstrate and prove agency's power and primacy for creating successful interactive experiences.

2 Embracing the Combinatorial Explosion

In this section we will focus on the quintessential requirements of interactive story, and along the way, prescribe an agenda for R&D in the field.

2.1 Agency Is Primary

To reiterate the key point from our comparison of interactive stories to videogames, for interactive stories to realize their full potential as a new narrative form, *agency* is the primary feature that must be offered players. While the term “interactive” encompasses more than agency [1], we believe agency to be the most critical and impactful facet of interactivity; again, we point to the success of high-agency videogames as empirical proof of this assertion.

Also important, it can be argued that the first wish that most players, developers and researchers originally feel when first encountering and considering interactive story, is the implicit promise to the player to be able to directly affect the plot of the story, taking it in whatever direction they wish.

If that wish is not always in the forefront of people’s minds today when they experience or consider interactive stories, we believe it is only because they have been conditioned by play experiences (or research experiments) that fall short of this ideal, causing them to suspend, forget about or give up on that dream.

2.2 Generation

The assumption that player agency is a feature of primary importance for fully-realized interactive stories would mean that players will need to be enabled to direct the plot of the story in many, many different directions. It follows that such an interactive story system must contain vast amounts of story content, to support this capability.

Any developer who has tried brute-force to write a branching story that can go in many different directions knows that the amount of content to be written exponentially explodes only moments after the beginning of the story. It is impossible to implement an interactive story of any meaningful breadth or depth in this way.

Yet, the amount of story potential in that combinatorial explosion is in fact what researchers and developers of interactive stories must embrace. It is effectively the target for what it means to give players true agency.

The only tractable way to achieve these results is through the real-time generation of story content from smaller building blocks. Researchers and developers need to build systems that have the potential to generate the nodes of an impossibly-large branching story graph, in response to the player’s direction. (At any one time, only the nodes needed at that moment need to be generated, not the entire story graph.)

There is a history of story generation research that we can look to, including such seminal projects such as Talespin, Universe and Minstrel. While the experimental results of each of these systems may seem modest, and to some perhaps even disappointing, they are in fact useful attempts to push in the necessary direction for achieving player agency. Unfortunately, few contemporary researchers and developers are explicitly working on story generation, perhaps discouraged by the results of the previous research, or by the difficult nature of the problem. This resistance must be overcome in order to advance the field.

2.3 Interface

To support player agency, along with story content generation, players will need expressive interfaces. By definition, enabling a player to direct the plot of an interactive story in many different directions requires the means to express those varied directions.

As mentioned earlier, games have only successfully offered players agency in the domain of action- and puzzle-oriented gameplay. Player “moves” in these domains can be achieved via spatial controllers and event triggers, such as joysticks and buttons. However, once players are directing the plots of stories about people, natural language and gesture interfaces will be required. Research and development of NLU interfaces will be just as critical for realizing the full potential of interactive stories as R&D for story content generation.

Even more than story generation, there is a substantial history of NLU research and development, as alluded to in the earlier discussion of speech recognition. Yet, much of that research has been in the domain of task-oriented interfaces, such as digital assistants. The domain of art and entertainment is distinct enough from task-oriented interactions that it warrants and requires its own research effort.

2.4 Connecting Generation and Interface

Chris Crawford suggested a conversational metaphor to describe the architecture of an interactive story system, where the system cycles through phases of listening, thinking, and speaking [2]. If generation is akin to speaking, and interface is akin to listening, then the two phases need to be connected with the ability to reason about the player’s input, to drive the story generation. Planning and drama management fall in this category. (Generation requires planning and reasoning as well, but at a point further along in the cycle.)

Drama management has gotten more attention in recent years from the research community than generation or interface, however it will be necessary to integrate it with generation and interface to fully develop it.

2.5 Terminology Problem

If the focus of R&D is centered on agency, the term “interactive storytelling” becomes a misnomer. The verb “telling” implies a preconceived notion of something to tell. Such a conceptual framework is antithetical to the notion of giving primary control to players to direct the interactive story as they play. As we work towards building interactive stories with true agency, a new term emphasizing agency and generativity becomes necessary. “Interactive storymaking” would be an example of a small deviation from the original term, if awkward and imperfect. (In this paper we avoid use of a verb altogether, and just refer to the product of the system, “interactive stories”.)

(Videogames have a similar crisis of terminology: the traditional meaning of the term “game” does not encompass all products currently assigned to it, particularly non-goal-oriented interactive entertainment products such as *The Sims* or virtual pets, as well as for more “serious”, less recreation-oriented works.)

Looking back at the birth of cinema again, we realize that the term “movies” had to be coined, shorthand for “moving pictures”.

3 Proceeding Forward

We conclude this editorial with some brief suggestions of how to frame future research and development of interactive stories.

3.1 Form Follows Function

It is not uncommon for researchers to develop a general architecture for some subsystem of an interactive story architecture, such as a generalized autonomous agent architecture, planner or drama manager. The drawback of this approach is that it is unclear what features of the subsystem will actually be useful and usable when attempting to build a working, playable interactive story. (The building of playable works, of course, is what we strongly believe must be the overarching goal of any R&D in the field; research should not be conducted for its own sake.)

We strongly recommend that researchers begin their projects with a specific concept of an interactive story in mind to target, and build an architecture to accomplish that target – even if they never intend to fully build the target interactive story itself. When focused on solving a particular, specific problem, the chances of provable success are greater, and perhaps less obviously, the chances are greater that a more generalized architecture can be extracted out of the research, because it is built to implement a grounded, working system.

3.2 Release Publicly Playable Projects

Because interactive story is an artform, the public ultimately must receive and critique the work. Researchers that never get the chance to fully test their results on players outside of the academic community will inherently be less informed and less connected to the true requirements of their pursuit.

With digital distribution of software so prevalent today, researchers and developers should plan from the beginning to eventually release playable prototypes to the public, to gather valuable feedback and critique. This matches nicely with the earlier suggestion to build systems that attempt to accomplish specific interactive story designs.

References

1. Murray, J.: *Hamlet on the Holodeck*. MIT Press, Cambridge (1998)
2. Crawford, C.: *Fundamentals of Interactivity* (1993), http://www.erasmatazz.com/library/JCGD_Volume_7/Fundamentals.html

Interactive Narrative, Plot Types, and Interpersonal Relations

Marie-Laure Ryan

University of Colorado, Boulder
marilaur@gmail.com

Abstract. The design of an interactive narrative begins with the choice of a type of story. In this paper I examine the potential of three kinds of plot for active user participation: the epic plot, which focuses on the struggle of the individual to survive in a hostile world, the dramatic plot, which deals with the evolution of a network of human relations, and the epistemic plot, which is propelled by the desire to solve a mystery. I distinguish two basic types of immersion—ludic and narrative, the latter subdivided into spatial, temporal and emotional variants, and I discuss the ability of the three kinds of plot to generate these various forms of immersion.

Keywords: Interactive narrative, design of interactive narrative, computer games, agency, user role, plot types, epic plot, epistemic plot, dramatic plot, interpersonal relations, ludic immersion, narrative immersion.

1 Introduction

The design of an interactive narrative begins with the choice of a type of story. In this paper I propose to examine the potential of three kinds of plot for interactive implementation: the epic plot, the epistemic plot, and the dramatic plot.

1.1 Three Types of Plot

Aristotle recognized two forms of narrative: the epic and the dramatic. Although the distinction was primary based on the mode of presentation—the epic representing events through verbal narration (*diegesis*) and the dramatic through an imitation of action (*mimesis*)—the two genres also differ from each other through the form and content of the plot: as Aristotle wrote, “one should not compose a tragedy out of a body of materials which would serve for an epic—by which I mean one that contains a multiplicity of stories” [1].

The epic plot is focused on the exploits of a solitary hero; it “preserves the memory of glorious deeds performed by superior beings” who “show their mettle in battles against human foes, monsters, or the powers of nature” [3]. Since every feat adds to the glory of the hero, the story can be endlessly expanded by adding new feats and new episodes. Epic narratives focus on physical actions, and the human relations that motivate the hero to act remain fairly simple. Take the archetypal plot of the fairy

tale, as described by Vladimir Propp [10]: a villain causes harm to a family, typically by kidnapping a princess; the hero is dispatched to repair the situation; after a certain number of tests he fulfills his mission by defeating the villain with the help of a donor, and he is rewarded for his actions with the hand of the princess. Throughout this plot, there is no evolution in personal relations: the hero is the faithful servant of the dispatcher; he remains opposed to the villain until the end (there is no reconciliation); nobody changes side during the fight between the two factions; when the hero marries the princess, they live happily ever after.

Dramatic narratives, by contrast, focus on evolving networks of human relations. Here are some examples of this evolution: in the beginning, *x* is allied with *y*; then *x* betrays *y* and sides with *z*, and in the end *x* and *y* are mortal enemies; in the beginning *x* has always been faithful to *y*; then *x* falls in love with *z*; *y* becomes jealous and kills *x* and *z*; in the beginning *x* and *y* are friends, then *y* insults *x* and they are enemies; but in the end *x* redeems himself, *y* pardons *x* and they are friends again. In a dramatic plot, the action is mental rather than physical: most of the events consist of acts of verbal communication between the characters; and when the characters perform physical actions, the significance of these actions resides in what they reveal about the mind of the agent and in how they affect interpersonal relations. Another difference from epic plots is that dramatic narratives present a closed pattern of exposition, complication, crisis and resolution (also known as the Freytag triangle) that defies expansion. The focus on interpersonal relations of the dramatic plot describes both the tragic and the comic genre¹.

In the nineteenth century, a third kind of narrative made its appearance: the epistemic narrative, driven by the desire to know. Its standard representative is the mystery story. The trademark of the epistemic plot is the superposition of two stories: one constituted by the events that took place in the past, and the other by the investigation that leads to their discovery. While the seed of the genre can be found in earlier forms of narrative—for instance, in Sophocles's *Oedipus Rex*—Walter Ong [9] attributes its emergence to the invention of print. The intellectual appeal of the mystery story lies in challenging the reader to find the solution before it is given out by the narrative; in order to do so, the reader needs to sort out the clues from the accidental facts, and to submit these clues to logical operations of deduction and induction. This mental activity would not be possible if the print medium did not give her the opportunity to parse the story at her own pace.

2 The Aesthetic Goals of Interactive Narrative

Janet Murray [7] proposes the Holodeck of the TV series *Star Trek* as the ideal model of interactive narrative. The Holodeck is a computer-generated, three-dimensional simulation of a fictional world. The user is invited to step into this world, to impersonate a character, and to interact through language and gestures with synthetic

¹ With the development of the novel, and even earlier, in *The Iliad*, the prototypical epic and dramatic plot were hybridized into open plots and action-centered stories that present much more complex interpersonal relations than the archetypal fairy tale.

agents. No matter what the user says or does, the synthetic agents respond coherently, and integrate the user's input into a narrative arc that sustains interest. The Holodeck, as a whole, may be a castle in the air, but this does not take anything away from the validity of its individual features as goals to pursue for researchers and developers of interactive narrative, even if it is often necessary to sacrifice one goal to another.

- Natural interface. Ideally, users should interact with computer-generated worlds in exactly the same way they interact with the real world: through language and the body. But for practical reasons, interactive narrative must often settle for less natural, but more efficient interfaces: the symbolic performance of physical actions through the manipulation of controls, and the selection of items from a language-based menu.
- Integration of user actions within the story. Just as, in real life, all of our actions contribute to our life story, in an interactive narrative, all of the user's actions should move the plot forward. This means that these actions should be more than a mean to unlock the next episode in a story told primarily through film clips and other non-interactive media. There should be a close thematic relation between the tasks offered to the user and the plot of the story.
- Frequent interaction. In real life, we interact with people and the world on a fairly constant basis, though there are moments when we are stripped of agency and forced to watch as spectators the events that determine our destiny. To reproduce this aspect of life, interactive narrative should make interactive moments the rule and passive moments the exception, rather than limiting agency to a few decision points separated by long stretches of passive watching.
- Dynamic creation of the story. Though interactive narrative will always rely on pre-scripted elements, the plot should be created as much as possible in the real-time of the user's interaction with the system. A reasonable number of different variations should emerge from different visits.

While the purpose of the preceding goals is to bring interactive narrative as close as possible to life, the next goal transcends the simulation of life by elevating it into art.

- Ability to create narrative immersion. By narrative immersion I mean an engagement of the imagination in the mental construction and contemplation of a storyworld. I contrast it with ludic immersion, which is a deep absorption in the performance of a task, comparable to the intensity with which a violinist performs a concerto. The relative importance of ludic and narrative immersion with respect to each other spells the difference between narrative games, in which the player pursues a specific goal associated with winning, and playable stories, in which the production of an aesthetically rewarding story is a goal in itself. Narrative immersion can take at least three forms [11]: spatial (a sense of place and pleasure taken in exploring the storyworld), temporal (a burning desire to know what will happen next) and emotional (affective reactions to the story and to the characters).

3 Interactivity and the Epic Plot

The epic plot has long been a favorite of game designers. We find it in shooters, in adventure games, and in the quests of MMORPGs. It is easy to see why it is so popular. The most common mode of interaction in computer games is through the keyboard, or through game pads. The range of actions that can be symbolically performed in real time through these controls is limited to the physical kind: players can move the body of their avatar, inspect or pick up objects by clicking on them, and trigger the behaviors encoded in these objects, such as firing weapons. The archetypal narrative pattern of the quest makes the most out of these limitations. The deeds of the hero are relatively easy to simulate through the game controls, the basic sequence of accomplishment-reward can be repeated endlessly, allowing the player to reach higher and higher levels in the game, the basic script of the quest lends itself to great variations in setting and in the nature of the tasks, and the solitary nature of the hero's quest makes interpersonal relations dispensable. Even when the player needs the assistance of system-created characters or of other players to perform the tasks (as happens in MMORPGs), he advances in the game on his own, and other characters are usually reduced to the fixed roles of either antagonists or helpers. Just as the epic genre highlights the physical deeds of the hero, the games based on this narrative pattern give players the opportunity to distinguish themselves by performing "exploits"—a term which, in hacker jargon, designates an extraordinarily daring and original way to accomplish a task, usually by taking advantage of a hole in the system. Another reason for the popularity of the epic pattern in video games lies in the graphic capabilities of computers. Epic narratives are basically travel stories that recount the adventures of the hero in a world full of danger. The ability of 3D graphic engines to adapt the display to the position of the player's virtual body makes them very efficient at simulating movement as an embodied experience, thereby creating spatial immersion.

4 Interactivity and the Epistemic Plot

The epistemic plot runs a close second to the epic plot in its compatibility with user interaction. It casts the player in the well-defined role of detective, it combines an authorially defined story—the events being investigated—with a variable story created in real time by the actions of the player; it takes advantage of the visual resources of digital systems by sending the player on a search for clues disseminated throughout the storyworld, and it is fully compatible with the types of action that can be easily performed by game controls: moving across the world, picking objects, examining them for clues, finding documents, and interrogating non-playing characters, ideally through a dialogue system but, more efficiently, through a menu of canned questions. An occasional fight against an antagonist can be thrown in to raise the adrenaline level of the player. The desire to know that drives the epistemic plot situates the immersivity of the genre on the temporal level. Temporal immersion includes three narrative effects [12]: curiosity, surprise and suspense. The first two dominate the mystery story, while the third is more typical of thrillers. When participation takes the form of spatial exploration and leads to unexpected discoveries, its motivation is

curiosity, and its reward is surprise. It is relatively easy for a game system to take away control from the player and to throw in events that suddenly solve the mystery in a totally unexpected way. Suspense however is much more resistant to interactivity, because it requires a long-range planning by the system and a top-down management of the player's expectations. Like epistemic curiosity, suspense is created by an intense desire to know, but while epistemic curiosity concerns events that already happened, suspense is focused on the future. People experience suspense when they can foresee two or more possible developments, and they are dying to find out which one of these paths the story will actualize. But when players can determine the path through their choice of actions, the uncertainty is lost. And if the system generates an accidental event to prevent the player from fully controlling the outcome of the story, the effect will be surprise rather than suspense.

5 Interactivity and the Dramatic Plot

The dramatic plot is the most difficult to implement because of its emphasis on the evolution of interpersonal relations. In the goal-oriented action of narrative games based on epic and epistemic plots, characters mostly matter to the player because of their capacity to help or hinder the achievement of tasks. Their relation to the player is fixed, and their relations among themselves are practically non-existent. Very rare are the cases in which the player regards NPCs as human beings rather than as means toward an end. According to Michael Nitsche [8], such a situation occurs in *Deus Ex* when the player must kill a formerly friendly NPC who has turned into a zombie in order to progress in the game. Some players develop such attachment to this characters that they experience extreme discomfort at treating her like an object that needs to be eliminated. The change in relation from friendly to inimical, and the moral ambiguity felt by the players—an ambiguity that temporarily distracts them from the game goals—represents a small step from an epic, strictly goal-oriented narrative game to a playable story with a dramatic plot.

The implementation of the dramatic plot raises countless problems. What will be the goals of the player and what kind of algorithm will it take to make these goals interact with the goals of the system-created characters? In an epic plot the player is given a goal by the dispatcher, and all of his efforts are geared toward the accomplishment of the mission. Similarly, in an epistemic plot, the player remains focused on the elucidation of the mystery until he finds the solution. But in a dramatic plot with evolving interpersonal relations, the goals of characters evolve together with their relations, and they must be constantly redefined. This requires of the system a much more powerful ability to simulate human reasoning than in epic and epistemic plots. Will players spontaneously adapt the goals and plans of their avatar to the current situation, or will the system tell them what to do through a NPC? Will NPCs be equipped with an AI sufficiently sophisticated to read each other's mind, as well as the mind of the player, and to adapt their beliefs, wishes, and plans to every change in situation? The emphasis of the dramatic plot on mental states requires from all the participants—whether human or system-created—an ability to form what is known in cognitive science as “theory of mind,” this is to say, representations of the thoughts of other characters. Does she love me or doesn't she? Is he trying to help or to deceive

me? Does she intent to keep her promise? Does she know that I know that she loves him? Can I believe what he says, or is he lying? These are the questions that propel and motivate the dramatic plot—whether tragic or comic—, the questions that the characters must ask in order to know how to respond to the actions of other characters. The creation of dynamic interpersonal relations between the player and the characters and between the characters themselves makes enormous, if not unrealistic demands on the AI that runs the system.

When—and if—all these issues are resolved, the question will remain of what kind of role should be given to the player, in order to make the visit to the storyworld into a truly pleasurable experience. I seriously doubt that people would enjoy emulating the heroes of tragedy and comedy by turning themselves into objects of pity or laughter. This suggests that the safest role for the player is that of a marginally involved observer or confidante whose interventions serves mainly as a stimulant that affects the behavior of the synthetic characters and triggers changes in their relations. Or, as Cavazza *et al.* suggest [2], the player could use his agency to change the world, in order to see how the characters will adapt to the new circumstances. As a peripheral character the user combines the roles of agent and spectator without assuming responsibility for the development of the story and without relinquishing the guidance of an author. As Glassner observes [4], most users do not really want to become improvising actors. If they are truly interested in creating their own scripts, they will be much better off participating in online virtual worlds where they will be able to interact with naturally intelligent agents.

5.1 Emotional Immersion in Interactive Narratives

The trademark of the dramatic plot lies in its ability to create an emotional type of immersion. In real life we experience two main types of emotions: those directed toward ourselves, and those directed toward other creatures through a vicarious experience known as empathy. Self-directed emotions concern our desires and the success of the actions through which we try to fulfill them. Even when these emotions involve feelings toward others, such as love and jealousy, the other is an object in a bi-polar relation determined by the desires of the experiencer. Not so with empathy: it is by mentally simulating the situation of others, by pretending to be them and imagining their desires as our own that we feel joy, pity, or sadness for them.

Passive media such as film, theater and novels have a unique power to generate emotions directed toward others. Aristotle paid tribute to this ability when he described the effect of tragedy as purification (*catharsis*) through terror and pity. By contrast, the emotions we experience while playing games—excitement, triumph, dejection, relief, frustration, relaxation, curiosity, and amusement [5]—are overwhelmingly self-directed ones, because they reflect our success and interest in playing the game. But their range is much smaller than the self-centered emotions of life: computer game players may fight to rescue a princess, and they may receive her hand in reward, but unlike the heroes of love stories, they are not motivated to act by romantic feelings.

While narrative games based on epic patterns deliberately sacrifice characters to action, playable stories based on interpersonal relations have only been able to create characters sufficiently lifelike to generate emotional reactions by limiting the player's

agency. A case in point is the interactive drama *Façade* [6]. The player develops intense feelings of dislike and contempt for Grace and Tripp, a professionally and socially successful couple whose seemingly perfect marriage turns out during the dialogue to be a mere façade that hides a deeply fractured personal relation. The player's visit with Grace and Tripp may act as a catalyst of the couple's hidden feelings, but she is limited to an observer role, and while her agency allows variations in the dialogue that exposes diverse facets of Grace and Tripp's personalities, the drama unfolds according to a relatively fixed script imposed top-down by the system. She may hate or despise Grace and Tripp, but unlike the player of a competitive game or online world, she does not entertain strong feelings for her avatar, such as caring for her character's personal relationship to Grace and Tripp. I certainly did not experience sadness over a lost friendship when the couple expelled me from their apartment at the end of the evening to sort out their problems between themselves.

6 Conclusion

With the epic quest structure of most video games, interactive media have mastered what could be the oldest form of narrative (or at least the oldest form of fictional narrative, for gossip must be older): the struggle of the individual against a hostile world. What remains to be conquered is the dramatic narrative, the plot form that knots together several destinies into a dynamic network of human relations and then disentangles them to let characters go their own way. Some steps in this direction have been taken with games interspersed with filmic clips and with interactive drama, but in both of these approaches, the involvement of the player remains peripheral: with film clips he relinquishes agency while the plot is being knotted, while with interactive drama (or at least with *Façade*, currently to my knowledge the only publicly released example of the genre), he is confined to an observer role. As the example from *Deus Ex* suggests, it may be relatively feasible to introduce system-created characters who inspire emotions to the player in a narrative game with an epic structure. Though emotional involvement conflicts with the pursuit of the game goals, these goals can always be temporarily suspended and resumed later. An online world could for instance be filled with interesting characters who tell moving stories to the player, momentarily distracting her from her current quest, or a friendly character could be turned into an enemy, causing a fleeting sense of betrayal in the player. It is much harder, and perhaps not desirable, to engage the player as a protagonist driven by strong self-centered desires in a playable story with a dramatic plot, because if the player identifies strongly with her avatar, the drama could become too much like life itself: an occasion for suffering and frustration, rather than a pleasurable experience. Finding a formula for an enjoyable combination of the self-centered emotions that come from our active engagement in a world or in a game with the other-centered emotions inspired by the fictional characters of narrative is the most daunting of the tasks that awaits the designers of narrative games and playable stories.

References

1. Aristotle: *Poetics*. Penguin, London (1996)
2. Cavazza, M., Charles, F., Mead, S.: Characters in Search of an Author: AI-based Virtual Storytelling. In: Balet, O., Subsol, G., Torguet, P. (eds.) *ICVS 2001*. LNCS, vol. 2197, pp. 145–154. Springer, Heidelberg (2001)
3. de Jong, I.: Epic. In: Herman, D., Ryan, M.-L., Jahn, M. (eds.) *The Routledge Encyclopedia of Narrative Theory*, Routledge, London (2005)
4. Zancanaro, M., Cappelletti, A., Signorini, C., Strapparava, C.: An Authoring Tool for Intelligent Educational Games. In: Balet, O., Subsol, G., Torguet, P. (eds.) *ICVS 2001*. LNCS, vol. 2197, pp. 61–68. Springer, Heidelberg (2001)
5. Lazzaro, N.: *Why We Play Games: Four Keys to More Emotions Without Story*, http://www.xeodesign.com/xeodesign_whyweplaygames.pdf
6. Mateas, M., Stern, A.: *Façade*, <http://www.interactivestory.net/download/>
7. Murray, J.: *Hamlet on the Holodeck. The Future of Narrative in Cyberspace*. Free Press, New York (1997)
8. Nitsche, M.: *Video Game Space*. MIT Press, Cambridge (2008)
9. Ong, W.: *Orality and Literacy: The Technologizing of the Word*. Methuen, London (1982)
10. Propp, V.: *Morphology of the Folk Tale*. University of Texas Press, Austin (1968)
11. Ryan, M.-L.: *Narrative as Virtual Reality: Immersion and Interactivity in Literature and Electronic Media*. Johns Hopkins University Press, Baltimore (2001)
12. Sternberg, M.: Telling in Time (II): Chronology, Teleology, Narrativity. *Poetics Today* 13, 463–541 (1992)

The IRIS Network of Excellence: Integrating Research in Interactive Storytelling

Marc Cavazza¹, Stéphane Donikian², Marc Christie², Ulrike Spierling³,
Nicolas Szilas⁴, Peter Vorderer⁵, Tilo Hartmann⁵, Christoph Klimmt⁵,
Elisabeth André⁶, Ronan Champagnat⁷, Paolo Petta⁸, and Patrick Olivier⁹

¹ School of Computing, University of Teesside, TS1 3BA Middlesbrough, United Kingdom

² IRISA/INRIA, Campus de Beaulieu, 35042 Rennes, France

³ FH Erfurt, University of Applied Sciences, Altonaer Str. 25, 99085 Erfurt, Germany

⁴ TECFA, FPSE, University of Geneva, CH 1211 Genève 4, Switzerland

⁵ Center for Advanced Media Research Amsterdam (CAMERA),

VU University Amsterdam

De Boelelaan 1081, 1081 HV Amsterdam, The Netherlands

⁶ Institute of Computer Science, University of Augsburg, 86135 Augsburg, Germany

⁷ University of La Rochelle, Pôle Sciences et Technologies, 17042 La Rochelle, France

⁸ Austrian Research Institute for Artificial Intelligence, Freyung

6/6, A-1010 Vienna, Austria

⁹ Culture Lab, School of Computing Science, Newcastle University, Newcastle Upon Tyne,
NE1 7RU, United Kingdom

m.o.cavazza@tees.ac.uk, stephane.donikian@irisa.fr,

marc.christie@irisa.fr, spierling@fh-erfurt.de,

Nicolas.Szilas@unige.ch, p.vorderer@fsw.vu.nl,

t.hartmann@fsw.vu.nl, klimmt@uni-mainz.de,

andre@informatik.uni-augsburg.de, ronan.champagnat@univ-lr.fr,

paolo.petta@ofai.at, p.l.olivier@ncl.ac.uk

Abstract. Interactive Storytelling is a major endeavour to develop new media which could offer a radically new user experience, with a potential to revolutionise digital entertainment. European research in Interactive Storytelling has played a leading role in the development of the field, and this creates a unique opportunity to strengthen its position even further by structuring collaboration between some of its main actors. IRIS (Integrating Research in Interactive Storytelling) aims at creating a virtual centre of excellence that will be able to progress the understanding of fundamental aspects of Interactive Storytelling and the development of corresponding technologies.

Keywords: Narrative Formalisms, Planning, Authoring Tools, Character Animation, Camera Control, Interactive Storytelling. Interactive Narrative.

1 Interactive Storytelling

Even though Interactive Storytelling (IS) can be traced back to early experiments with interactive cinema, the modern concept of IS has been formulated during the 1990s, through the work of Davenport [9], Nakatsu and Tosa [22] and the OZ project [2].

Yet, research in the field only developed over the last decade, in particular with the emergence of Artificial Intelligence techniques for story generation. The IRIS (Integrating Research in Interactive Storytelling) Network of Excellence is a new EC-funded initiative (under FP7's Intelligent Content and Semantics activities) whose work should commence exactly 10 years after the 1999 AAAI Fall Symposium on Interactive Narrative [18], which can be considered as a milestone in the emergence of the discipline. This Network brings together research institutions and individuals which have contributed to the development of this research topic. It aims at advancing the discipline by providing a more integrated approach to IS both in terms of theoretical foundations and in terms of technology development. In a sense, it is a natural follower of previous projects such as INSCAPE [38], which has played an important role in the development of this Conference, through its predecessors, the ICVS (2001-2007) and TIDSE (2003-2006) conferences. However, its specificity as a NoE is to advance knowledge in the discipline itself, rather than promote a specific approach or develop a specific technology.

2 State-of-the-Art and Open Challenges

Rapid progress over the past ten years has materialised in the development of several Research prototypes [1] [4] [6] [10] [16] [17] [20] [23] [24] [28] [30] [33] [37]. A number of key concepts have emerged as central to IS, amongst which narrative control [32], duality between character and plot [19], the potential of Planning techniques for action generation [34] [36]. In addition, there has been convergence in the recognition of open problems and challenges for the development of IS systems:

- **Relations between Action generation and Staging.** The various “Story Engines” developed in IS have mostly focussed on action generation using AI techniques (Planning being the dominant one). The development of principled relations to the actual staging of such actions (their visual presentation [8] [15]) has been somehow neglected.
- **Relations between Narrative Formalisms and AI techniques.** References to narratology abound in IS research (see [5] for a review), but do not always translate into the actual implementation in particular at the level of AI formalisms.
- **Authoring.** Authoring is now considered a central problem for IS [13] [21] [25] [26] [27] [38], and the development of authoring tools a major condition for the success of IS technologies. The challenge for authoring is to evolve from proprietary solutions attached to specific IS systems.
- **User Interaction for Interactive Storytelling.** IS requires the integration of story generation and natural user interaction mechanisms. Many systems incorporate natural language processing or dialogue in search of natural interaction [4] [6] [20] [28] [33]. However, most of the effort remains dedicated to story generation, perhaps in the hope that IS could directly benefit from progress in multimodal interfaces, and some more situated interaction research is certainly needed.
- **Evaluation Methodologies for Interactive Storytelling.** Whilst there have been some considerations on measuring technical performance and scalability [7] [20]

(also for authoring [3]) and on high-level properties of the narrative such as suspense [35], there are so far no established methodologies to evaluate the “quality” of an Interactive Narrative.

Of these challenges, we shall briefly discuss two, which, in our view, have not received sufficient attention to date.

3 Computational Narratology

The theoretical foundations of IS are comparatively less advanced, and most of the above concepts have emerged as empirical findings in the development of the above prototypes and associated experiments. Despite the identification of Narratology [29] and Narrative formalisms [5] [11] as a potential theoretical background, their role in IS is largely underspecified if compared to the role of Linguistics for Computational Linguistics and Natural Language Processing; in other words, whether IS can be grounded in Computational Narratology is still a matter to be explored. Narrative Formalisms can influence all aspects of IS, from story generation to story presentations. They can influence the definition of AI formalisms for story generation, whether these are based on planning or not. Understanding narrative formalisms also helps story presentation through aspects such as camera placement: to a large extent, film idioms [8] popular in camera control correspond to empirical narrative formalisms and could benefit from a more fundamental analysis. Ultimately narrative formalisms can support unified representations for narrative actions which are also important for authoring and even the definition of software architectures (for instance providing principled interfaces between action and virtual actors’ animation).

4 Evaluating Interactive Storytelling

While there is strong tradition in the social sciences and the humanities to investigate readers’ (and, more recently, viewers’) experience of and response to conventional linear stories such as novels, short stories, or movies, user-centred research on IS is not a well-established field. A few research units worldwide run studies on the systematic measurement of user responses to video games. Some small scale pilot studies have been published on user reactions to dedicated IS environments, mostly for the purpose of specific system optimisation or the initial demonstration that a given application ‘can work’ (e.g., [1] [14]). However, the methodologies applied (also in video game research) are highly diverse, and there is not much practical experience with effective measures for system evaluation in IS available. While the first evaluation issue relates to general system usability, target audiences and the question which kind of people are likely to select and enjoy a given IS environment, the second issue reflects the ‘entertainment value’ or enjoyment of the experience delivered by the system. Enjoyment is a highly complex experiential state with a variety of manifestations (e.g., exhilaration, suspense, pride) and numerous determinants attached to both the system delivering the experience and the person confronted with the system [31]. To evaluate whether a given IS system achieves a satisfying level of enjoyment in

users from defined audiences, it is thus necessary to specify in advance the kind of experiential quality a system shall deliver. Consequently, experimental exposure studies with control groups can apply appropriate enjoyment measures and reveal the entertainment-related capacity of the system.

5 Conclusions

Besides the joint research that will take place within IRIS, the Network plans to develop activities to the benefit of the larger IS Research community. It will endeavour to disseminate theoretical results, technological components and data, for instance evaluation methodologies, comparative formalisations of plots, and various corpora (for instance of interaction data).

References

1. Aylett, R.S., Louchart, S., Dias, J., Paiva, A., Vala, M.: FearNot! - An experiment in emergent narrative. In: Panayiotopoulos, T., Gratch, J., Aylett, R.S., Ballin, D., Olivier, P., Rist, T., et al. (eds.) IVA 2005. LNCS, vol. 3661, pp. 305–316. Springer, Heidelberg (2005)
2. Bates, J.: The role of emotion in believable agents. *Communications of the ACM* 37(7), 122–125 (1994)
3. Carbonaro, M., Cutumisu, M., McNaughton, M., Onuczko, C., Roy, T., Schaeffer, J., Szafron, D., Gillis, S., Kratchmer, S.: Interactive Story Writing in the Classroom: Using Computer Games. In: International DiGRA Conference, June 2005, pp. 323–338 (2005)
4. Cavazza, M., Charles, F., Mead, S.J.: Character-based Interactive Storytelling. *IEEE Intelligent Systems* 17(4), 17–24 (2002)
5. Cavazza, M., Pizzi, D.: Narratology for Interactive Storytelling: a Critical Introduction. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 72–83. Springer, Heidelberg (2006)
6. Cavazza, M., Lugin, J.-L., Pizzi, D., Charles, F.: Madame Bovary on the Holodeck: immersive interactive storytelling. In: ACM Multimedia 2007, pp. 651–660 (2007)
7. Charles, F., Cavazza, M.: Exploring the Scalability of Character-based Storytelling. In: Third ACM Joint Conference on Autonomous Agents and Multi-Agent Systems, New York, USA, pp. 872–879 (2004)
8. Christie, M., Olivier, P.: Automatic Camera Control. State of the Art Report. In: Eurographics, Vienna, Austria (2006)
9. Davenport, G.: Seeking Dynamic, Adaptive Story Environments. *IEEE MultiMedia* 1(3), 9–13 (1994)
10. Donikian, S., Portugal, J.-N.: Writing Interactive Fiction Scenarii with DraMachina. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 101–112. Springer, Heidelberg (2004)
11. Grasbon, D., Braun, N.: A Morphological Approach to Interactive Storytelling. In: Proceedings of Cast 2001, Living in Mixed Realities, Sankt Augustin, Germany, pp. 337–340 (2001)
12. Gratch, J.: Why you should buy an emotional planner. In: Proceedings of the Autonomous Agents 1999 Workshop on Emotion-based Agent Architectures, EBAA 1999 (1999)

13. Iurgel, I.: Cyranus – an authoring tool for interactive edutainment applications. In: Pan, Z., et al. (eds.) *Edutainment 2006*. LNCS, vol. 3942, pp. 577–580. Springer, Heidelberg (2006)
14. Kelso, M.T., Weyhrauch, P., Bates, J.: Dramatic presence. *Presence: Teleoperators and Virtual Environments* 2(1), 1–15 (1993)
15. Lamarche, F., Donikian, S.: Automatic orchestration of behaviours through the management of resources and priority levels. In: *First Joint Conference on Autonomous Agents and Multi-Agents Systems (AAMAS 2002)*, pp. 1309–1316 (2002)
16. Machado, I., Paiva, A., Brna, P.: Real Characters in Virtual Stories (Promoting Interactive Story-Creation Activities). In: Balet, O., Subsol, G., Torguet, P. (eds.) *ICVS 2001*. LNCS, vol. 2197, pp. 127–134. Springer, Heidelberg (2001)
17. Magerko, B., Laird, J.E., Assanie, M., Kerfoot, A., Stokes, D.: AI Characters and Directors for Interactive Computer Games. In: *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI 2004)*, pp. 877–883 (2004)
18. Mateas, M., Sengers, P.: *AAAI Fall Symposium on Narrative Intelligence*. AAAI Technical Report FS-99-01 (1999)
19. Mateas, M., Stern, A.: Towards Integrating Plot and Character for Interactive Drama. In: *AAAI Fall Symposium on Socially Intelligent Agents: The Human in the Loop* (2000)
20. Mateas, M., Stern, A.: A Behavior Language for Story-based Believable Agents. *IEEE Intelligent Systems* 17(4) (2002)
21. Medler, B., Magerko, B.: Scribe: A Tool for Authoring Event Driven Interactive Drama. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) *TIDSE 2006*. LNCS, vol. 4326, pp. 139–150. Springer, Heidelberg (2006)
22. Nakatsu, R., Tosa, N.: Interactive Movies. In: Furht, B. (ed.) *Handbook of Internet and Multimedia – Systems and applications*. CRC Press and IEEE Press (1999)
23. Rank, S., Petta, P.: From ActAffAct to BehBehBeh: Increasing Affective Detail in a Story-World. In: Cavazza, M., Donikian, S. (eds.) *ICVS-VirtStory 2007*. LNCS, vol. 4871, pp. 206–209. Springer, Heidelberg (2007)
24. Riedl, M.O., Young, R.M.: An Intent-Driven Planner for Multi-Agent Story Generation. In: *Third ACM Joint Conference on Autonomous Agents and Multi-Agent Systems*, New York, USA, pp. 186–193 (2004)
25. Skorupski, J., Jayapalan, L., Marquez, S., Mateas, M.: Wide Ruled: A Friendly Interface to Author-Goal Based Story Generation. In: Cavazza, M., Donikian, S. (eds.) *ICVS-VirtStory 2007*. LNCS, vol. 4871, pp. 26–37. Springer, Heidelberg (2007)
26. Spierling, U., Weiß, S.A., Müller, W.: Towards Accessible Authoring Tools for Interactive Storytelling. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) *TIDSE 2006*. LNCS, vol. 4326, pp. 169–180. Springer, Heidelberg (2006)
27. Spierling, U.: Adding Aspects of “Implicit Creation” to the Authoring Process in Interactive Storytelling. In: Cavazza, M., Donikian, S. (eds.) *ICVS-VirtStory 2007*. LNCS, vol. 4871, pp. 13–25. Springer, Heidelberg (2007)
28. Swartout, W., Hill, R., Gratch, J., Johnson, W.L., Kyriakakis, C., LaBore, C., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Morie, J., Rickel, J., Thiebaut, M., Tuch, L., Whitney, R., Douglas, J.: Toward the Holodeck: Integrating Graphics, Sound, Character and Story. In: *Proceedings of the Autonomous Agents 2001 Conference* (2001)
29. Szilas, N.: Interactive Drama on Computer: Beyond Linear Narrative. In: *Proc. AAAI Fall Symposium on Narrative Intelligence* (North Falmouth MA). AAAI Press, Menlo Park (1999)
30. Szilas, N.: IDtension: a narrative engine for Interactive Drama. In: Göbel, S., et al. (eds.) *Proc. TIDSE 2003*, pp. 187–203. Fraunhofer IRB Verlag (2003)

31. Vorderer, P., Klimmt, C., Ritterfeld, U.: Enjoyment: At the heart of media entertainment. *Communication Theory* 14(4), 388–408 (2004)
32. Weyhrauch, P.: Guiding Interactive Drama, Ph.D. Dissertation, Tech report CMUCS-97-109, Carnegie Mellon University (1997)
33. Weiß, S., Müller, W., Spierling, U., Steimle, F.: Scenejo – An Interactive Storytelling Platform. In: Subsol, G. (ed.) *ICVS-VirtStory 2005*. LNCS, vol. 3805, pp. 77–80. Springer, Heidelberg (2005)
34. Young, R.M.: Notes on the use of Plan Structures in the Creation of Interactive Plot. In: *AAAI Fall Symposium on Narrative Intelligence*. AAAI Press, Menlo Park (1999)
35. Young, R.M.: Cognitive and computational models of suspense: Towards the automatic creation of suspense in interactive narratives. In: *Interactive Frictions: The Conference Produced at the Pressure Point Between Theory and Practice: An International Conference on Interactive Narrative*, University of Southern California, Los Angeles, CA (1999)
36. Young, R.M.: Creating Interactive Narrative Structures: The Potential for AI Approaches. In: *AAAI Spring Symposium in Artificial Intelligence and Entertainment*. AAAI Press, Menlo Park (2000)
37. Young, R.M.: An Overview of the Mimesis Architecture: Integrating Intelligent Narrative Control into an Existing Gaming Environment. In: *The Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, Stanford, CA (March 2001)
38. Zagalo, N., Göbel, S., Torres, A., Malkewitz, R.: INSCAPE: Emotion Expression and Experience in an Authoring Environment. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) *TIDSE 2006*. LNCS, vol. 4326, pp. 219–230. Springer, Heidelberg (2006)

Mobile Urban Drama – Setting the Stage with Location Based Technologies

Frank Allan Hansen, Karen Johanne Kortbek, and Kaj Grønbæk

Center for Interactive Spaces,
Department of Computer Science, University of Aarhus,
Aabogade 34, DK-8200, Aarhus N, Denmark
{fah,kortbek,kgronbak}@daimi.au.dk

Abstract. This paper introduces the novel concept of location-based *Mobile Urban Dramas*. In a Mobile Urban Drama the user become the main character in a play where actors' voices appear in the mobile phone headset linked to the physical setting in the city as the stage for the drama. The paper describes the dramaturgical concept and introduces a software framework supporting drama writers in developing such Mobile Urban Dramas. Experiences with use of the framework are discussed with successful examples of real dramas that have been developed and performed by a Danish theatre group, Katapult.

Keywords: Mobile drama, location-based technology, Semacodes, interactive audio, sensor/actuator based interaction, urban computing.

1 Introduction

Anna: “If you can do this, you will see me one more time. However, if you fail, you will end up talking to me for the rest of your life.”

Christian: “You can't sit there and threaten me.”

Anna: “My condolences. Goodbye.”

Christian: “Anna?... Anna?!”

(Pause)

Narrator: “You must go and take a picture of tag three. It is hanging at the entrance of the pub.”

(The end of scene 2 in the play *GAMA*).



Fig. 1. Tags anchor scenes in the city

Mobile Urban Dramas are interactive audio plays that let the spectator (hereafter called ‘user’) be the main character in a drama where the real urban environment becomes the scenography. In the play, users are equipped with mobile phones, headsets and maps, and experience a drama in the streets of a city, where they trigger different scenes of the play through location-based technology such as GPS or Semacode¹ based tags. Furthermore, they receive what appears to be SMS messages and phone calls, and encounter real life actors as part of the play that takes place along a pre-specified, possibly branching path through a cityscape or landscape. The

¹ <http://semacode.com/>

Danish theatre company Katapult [10] has developed three such plays using our technology. Katapult has developed these plays based on what they call an "AudioMove" concept. The AudioMove concept is structured as a single user experience, while the Mobile Urban Drama concept presented in this paper can be both single and multi user experiences. These plays are to our knowledge among the world's first applications of mobile phones for location-based theatre/drama.

Only few previous examples of using phones for drama exist. The most prominent example launched in 2005 by the German theatre collective Rimini Protokoll, as the world's first mobile phone theatre. The play was entitled "Call Cutta" [15] and took place in Kolkata, Calcutta, India. The play transforms the city into a stage and the people in the audience into main characters of the play. When the users buy a ticket at the box office, they may get the key for a room and a sketch of how to get there. It may be an arbitrary room in the urban space, e.g. an office, or an apartment. When entering the room they find a phone ringing, and the play then goes on with a call center person and the users as the main actors of the play, and the plot develops through the phone conversation. The room is the scenography; TVs or computers may turn on and integrate in the play. The play is today (2008) called "Call Cutta in a Box" and can be experienced in several cities around the world.

Rimini Protokoll are also known for a different mobile play "Cargo Sofia" [4], where the scenes take place in a Bulgarian truck driving in the streets of Dublin. Here the users are loaded into the container of a cargo truck, and they experience the drive through a video projection with audio on the container wall coming from the cab showing small film clips mixed with dialogue between the two drivers (actors). The users experience projections on the window, audio scenes in earbuds, and "meet" actors in the street when the truck appears at specific locations.

A final example of a mobile drama is the mobile TV project Song fuer C [19], which is a mobile detective story developed by Marc Weis and Martin DeMattia, Vodafone R&D, and the Hochschule für Gestaltung und Kunst Zürich (HGKZ). The user experiences video scenes on a Sony PSP or a mobile phone, while being guided around in some urban area where real actors may appear at certain locations.

For traditional computers, dramas that make the user into the main character are well known. One of the typical examples is Façade [12, 13]. It is a real-time, first-person dramatic world in which the player, visiting a married couple in their apartment, becomes entangled in the conflict dissolution of their marriage. The Façade interactive drama integrates real-time autonomous agents, and AI is applied [6] to manage the behavior of agents actors, thus the user never experience real or recorded human characters in this type of play.

This paper discusses the new genre of Mobile Urban Drama. The paper is structured as follows: Section 2 introduces the location based framework for managing the plot graph and the current implementation. Section 3 discusses three example plays and the initial experiences. Section 4 compares our concept to related work. Section 5 concludes the paper.

2 Setting the Stage with Location Based Technology

This section presents our conceptual framework for Mobile Urban Drama. The framework draws on elements from interactive storytelling [12, 17, 20], hyperfiction

[3], adaptive hypertext [5], and user modeling [11] in order to create a rich model for storytelling. We also discuss how to realize the framework by applying technologies that allow use of real physical locations as scenography for the dramas. The novelty of the work presented here, is the combination of classical techniques such as plot graphs for representing the narrative flow, user models for representing the user interaction and environment models for representing the physical environment and physical interaction through various input and output systems. The work thus builds on concepts from interactive storytelling [16, 17], urban storytelling [2, 14], and augmented reality fiction [18], but generalizes and integrates these concepts in order to create a tighter coupling between the urban scenography and the narrative and allow the interaction to be mediated by a combination of different pervasive technologies.

2.1 A Conceptual Framework for Mobile Urban Drama

In the conceptual framework, dramas are constructed as a series of events that form the story. These events can be depicted as a plot graph inspired by Ryan [17] as illustrated in Figure 2.




Fig. 2. A plot graph illustrating the basic elements of the framework. The plot is represented as a directed graph with guards that control possible paths. The guards specify node names (e.g. 1), sequences (e.g. 3,4), negations (e.g. !2), and / or values between these (e.g. 1 | 2).

The plot graph is a directed graph with nodes corresponding to events and edges specifying the valid ordering of the events. The plot graph must contain one or more start (S) and end events (E), where the story can begin and end, and any number of edges between the event nodes. The drama can be constructed as anything from a linear path through the nodes without any branches to a highly connected graph where the user freely can choose the next branch at each node. The graph can thus easily represent various interactive story architectures [17], e.g. trees that represent multiple variants of a story over time, flowcharts that are suitable for stories that contain episodes or levels such as games, and mazes, which represents hypertext like structures, where the plot is constructed from the users' actions at every node.

The basic plot graph only allows the author to specify events of the drama and their ordering. The author has no control of how the user actual follows the events. If the graph e.g. contains a loop, the user may experience the same event more than once and may in fact choose to follow the loop-sequences infinitely. To give the author more control of the plot, the framework includes two types of models, a *user model* and an *environment model*, which can be used in conjunction with the plot graph.

The user model is an overlay model [5] on the plot graph and describes the user's knowledge of the story. The model encodes the events the user has experienced, the ordering of the events and the choices taken at different branches in the story. Based on the user's current user model, the author can add *guards* (conditions) to the edges in the plot graph that controls whether or not that specific path will be available to the user. In this way the user's actions can have consequences at a later state in the story and the author can control the users possibilities e.g. based on the persons, the user has met or the places she has visited in the virtual world. As an example, consider Figure 2 above: the user may follow different paths in order to get to the event labeled 4. If the user has visited state 1 or 2, then she has the possibility to go through state 5 before ending the story, but if she has not experienced state 2 or has gone through state 3 and 4, the author sends her straight to the end in state E. In, this example the two paths are mutually exclusive, so based on the user's previous actions in start node S and node 1, the author controls the path the user must follow in the end.

The second model, the environment model, has a similar function as the user model, but is not directly controlled by the user's actions. The environment model is used to describe the physical environment of the play, e.g. at which physical locations events occur, what technologies are used to trigger the events, and so forth. The environment model also captures the global state of the play. For instance, if the drama supports multiple users participating in a play at the same time, some of their actions may be fed back to the environment model, and the choice of one user may influence the interaction space of the others. In Figure 2, an environment model guard is placed between state 1 and state 2. If one user has chosen to follow path from state 1 to state 2, the author could use the guard to make it less likely that this path would be available to other users at a later time of the play.

2.2 Technical Realisations of the Conceptual Framework

The conceptual framework discussed above only describes the interactive elements at the story level, but not how the drama is actually realized and integrated with real world scenography. However, the mapping between the conceptual model and the actual implementation is rather simple in that each story event (plot graph node), is triggered by a real event in the implementation. Events are triggered by some input from the user or a sensory system and the story event is provided by some kind of output to the user, as illustrated in Figure 3.

Inputs are used to trigger story events in the drama. Inputs may be anything from manual entry by the user, e.g. pressing a button on a mobile phone, to more automated input from a sensor, e.g. Semacodes, RFID, GPS, Bluetooth, or WiFi positioning systems. In terms of our system, it makes no difference how the input is generated, as long as it can be mapped to unique story events. However, the input technology will influence the interaction between the user and the real world scenography, and thus may impact the overall experience. Some technologies such as Semacode tags or RFID tags require the user to actively find and scan the tag either with the mobile phone's camera or an RFID scanner and the user must thus explicitly trigger the input.

Semacode tags are 2D patterns that encode a piece of text, typically an ID or a URL and can be used by software for camera phones. Such 2D barcodes can be used as digital-physical links. They have the advantage that they clearly visualize the link

in the physical world, and when decoded, provide the link address directly (i.e. the URL). Semacodes have been used in a number of pervasive applications [7]. Other technologies like GPS can be used to support implicit interaction, as the system can register when a user enters a certain location and use this to generate input events [9]. In this case the user will not have to explicitly generate the event, but simply move about in the physical world as part of the play.




Fig. 3. Applying technology to the plot graph

The output from the system can also use various technologies to optimize the experience. If mobile phones are used as interaction devices in the mobile drama, the phone can play audio and display text, graphics, and film clips on the screen and thus function as a versatile output device. However, the physical environment may also be used to present output. Displays in the user’s vicinity can show interactive graphics and (directional) speakers can provide audio or ambient effects. More permanent installations may also include 4D theater effects, such as wind, mist, or motion [1].




Fig. 4. Technical implementation of *CORRIDOR*. The plot graph is augmented with technologies that provide input, output, and timed events at each state.

In our implementations of the conceptual model, we have found timed output events to be very useful. Instead of relying solely on a tight input-output scheme, we have introduced timers, which allow us to suspend the event to a later time and thereby give the author more control of the output. Fig. 4 illustrates the implementation of the *CORRIDOR* drama (discussed in more detail in Section 5). The primary input technology in the setup was Semacode tags placed in the streets of Aarhus and the primary output was audio played through the user’s headset.

However, a number of the Semacode tags also triggered a timer that would send a text message to the user's mobile phone at a later time. At the final state E, the user would also receive a timed phone call and receive audio from speakers in the environment (which in this case were placed on an actual theater stage). The figure also illustrates how real actors were part of the *CORRIDOR* drama. The users would meet the actors while walking from one location to another, however, the actors did not provide input or output in a system sense, but interacted directly with the user.

The final aspect of an implementation that needs consideration is how to bridge the drama with the real world scenography. Fortunately, each of the input technologies discussed above make it relatively easy to anchor the plot to real physical locations. A GPS sensor attached to the mobile phone, for instance, can provide precise physical coordinates for a location. The story can then be written to take the physical location into account, e.g. by telling the user to "walk to the other end of the market square" when the system has registered that the user has entered the GPS-coordinate for the square. 2D barcodes such as Semacodes cannot provide such a global position, but since the tags are placed at physical locations, it is still possible to map the story events to those locations, i.e., the GPS example may be implemented simply by placing a Semacode close to the market square providing the same message as before.

Finally, it should be noted that there is not necessarily a one-to-one correspondence between story events and physical locations. Indeed, the author may often wish to create a branch between different output events at a given physical location in order to provide differentiated output that depends on the state of the user model at that particular location. However, it is of course also possible to create a direct mapping, which is the case in the *CORRIDOR* implementation depicted in Figure 4, where each state refers to a particular physical location and incorporates this into the story.

2.3 Implementation Used for the Three Dramas

Section 2.2 discussed how the conceptual framework can be realized by utilizing a number of alternative technologies. Here we discuss the technical implementations we have made so far. The applications presented here are built on a common urban computing infrastructure that we have developed for a variety of applications [8].

The applications rely on Semacode tags as input to anchor the story fragments to physical locations. Audio, video, and text is used as output on the mobile phones.

The software running the dramas was implemented on the Java JME platform and deployed on SonyEricsson K800i mobile phones, which include good camera support for scanning the tags and the possibility for using normal stereo headphones, for quality audio output. As the dramas included hundreds of megabytes of audio files, these files were not downloaded from the Web over the mobile Internet, but preloaded on a number of mobile phones that users would borrow when buying tickets.

Each Semacode tag encoded a unique URL and since the application used a fixed number of tags, each encoded URL was linked to a specific audio file in the application's environment model. When a tag was scanned it would trigger the corresponding audio file or timer event. However, the URLs also pointed to the Webpage of theatre Katapult and if scanned by a third-party Semacode application, they would redirect the phone's browser to an information page describing the audio theater and explaining how to buy tickets.

3 Example Plays and Initial Experiences

To date, the framework has been utilized in three Mobile Urban Dramas, “*CORRIDOR*”, “*GAMA – ON THE TRAIL OF UNKNOWN LAND*” and the outdoor school project “*HASLEINTERACTIVE*”, which will be launched this summer. The Danish theatre, Katapult [10], has developed all of the plays, which are described in the following.

3.1 Corridor

The first drama “*CORRIDOR*” was developed for the Aarhus Festival in 2007. The amount of scenes (tags) to be scanned was sixteen, and the duration of the play was approx. 1 hour and 40 min. The user is staged as a female journalist who is assigned to make a cover story on a secret PR manager due the next day. However, the PR manager keeps changing location, and moreover, a personal matter becomes urgent. As the plot progresses, the stress level increases and the journalists’ officious boss keeps phoning her. By walking in the streets of Aarhus, the user finds herself in a dramatised reality, where everyone could be a possible character in the play. Four actors and two actresses take part in the play, and moreover, the user receives events on the phone, that simulates SMS messages and calls. The staging of the Mobile Urban Drama is essential for the experience of the play and to whether or not the users can relate to the story, and accept the premise of being the main character. After having experienced the play, several users stated that they actually felt stressed out.

A total of more than 100 people experienced the play during the seven days it was staged last summer. Thirty-three of these completed an informal questionnaire consisting of open questions formulated by the theatre. On the question, “What is your overall experience of the play?” 94% answered, that they liked it and found it very exciting, and an interesting manner in which to experience the city.

On the question, “Was it difficult to navigate around the city?” 21% noted that the construction work in a particular part of the city made it more difficult, however, not impossible, to navigate. Furthermore, 13% would have preferred the ability to have the description repeated of which way to go; especially those users who did not know the city beforehand. However, most commented, that they found their way effortless. One commented on the map and another user on the ability to call for help on the service phone number as helpful features. On the question “Did you find it difficult to make use of the mobile phone – or to scan the tags?” 85% answered “No”. However, 12% had to scan the tags more than once, and one user stated, that being stressed out with somewhat shaking hands made it more difficult to scan the tags.

In the last question, users were asked if they had any suggestions for improvement. 33% would have liked an increased sound volume as the noisy environment could result in them not being able to hear the play. 18% would have liked to meet more actors in the play, or a larger degree of interaction. Finally, 12% suggested a larger degree of personal influence on the play – e.g. making choices. These results had an impact on the work on the second project, *GAMA*.

3.2 GAMA

The second project “*GAMA – ON THE TRAIL OF UNKNOWN LAND*” is played in the city of Horsens, and will play throughout the summer of 2008. The play consists of 14 scenes (tags) and the duration is 1 h. and 20 m. The story is about a medical student, Christian, who is on his way to his sisters funeral. On his way to church, he receives a mysterious phone call from a woman who claims to have a message from his sister, Johanne. The woman claims to be Anna Christina Bering, the long-dead wife of a famous explorer, Vitus Bering, who was descended from Horsens. The user is the main character in a psychological treasure hunt; and he is provided with historical facts of the city through his quest for messages from his late sister.

The overriding conceptual structure is the same in *CORRIDOR* and *GAMA*; both plays follow the classical Hollywood model, which is a linear course where the drama increases until a climax at the end of the plot. In the *GAMA* play, the user is more interactive than in *CORRIDOR*, as the user needs to find clues, e.g. old letters in drawers and other objects in his quest. As opposed to *CORRIDOR*, in *GAMA* several scenes take place indoor, and thus making better use of the surrounding environment than in *CORRIDOR*. Another considerable difference in the two urban dramas is the fact that in *GAMA* the locations were determined before the drama was written. This meant that the individual audio sequences were written accordingly to the approximate distance to the next location, which made less pauses in the storyline, and thus, the long music sequences that were utilized in *CORRIDOR* unnecessary. The duration of the longest audio clip in *GAMA* was eleven minutes, which has proven not to be too long – even to those users that walk quickly. Furthermore, the play could be made more place-specific when having determined the physical locations before the writing process. This makes the user experience more immersive, however, more difficult for the writers to compose having more restrictions than when dealing with classic theatrical dramas.

General experiences and lessons learned on urban dramas include not making the course through the city too complex. It should not be too difficult to find the way to the next tag as this would cause an increased focus on the map when navigating through the city, and less focus on the drama and the city itself. Furthermore, the tags should be carefully placed and protected.



Fig. 5. The user photographs tags as he walks in the city, which enables him to listen to the audio drama

3.3 HasleInteractive – Outdoor School

HASLEINTERACTIVE is a play being launched in September 2008. Unlike *CORRIDOR* and *GAMA*, it is developed for educational purposes for public schools; more specifically, sixth graders. Moreover, the narrative consists of both individual and social courses in the same play, as the audio sequences are supplemented by group work on tasks with close ties to the narrative. The duration of the play will be approx. three hours.

The narrative is developed for the hilly landscape of Hasle in the outskirts of Aarhus. The plotline is set in the year 2022 when large scenic areas of nature are dying out. Two scientists need help from children - only they can reveal the reason for the decaying nature. Adults cannot enter infected areas without being lethally drained of energy. The scientists do not understand the mysterious codes and events in the area, thus they send in the children, whose mission is to analyse, decipher mysterious codes, report back and find evidence from the biotopes - before it is too late.




Fig. 6. An episodic plot graph from the *HASLEINTERACTIVE* drama

The aim for their learning is to gain an understanding of the correlation between humans and nature. In doing this, each group of pupils are to solve four assignments in the subjects mathematics and natural sciences at four checkpoints – two located in different lake biotopes, and two in a young forest and an old forest, respectively (number 1-4 in Figure 6). When the pupils arrive at a checkpoint, they scan a tag and receive assignments developed in collaboration with teachers. The amount of tags in the play is six when counting in the tags at the beginning and at the end of the play. However, between the checkpoints are “coordination nodes” with mysterious signs on the trees providing the pupils with more abstract info related to the narrative (nodes a-f in Figure 6). These nodes tie the narrative by providing a relationship between individual sensuous audio drama experiences and social technical experiences.

Figure 6 shows a plot graph from the *HASLEINTERACTIVE* drama. The guards limit the user’s options and there are only four possible plots (given by the choice in node a). One of the reasons for creating the structure in Figure 6 is to allow more groups of users to use the same (physical) scenography at a time. Four different groups of three-four pupils may experience the play without interfering with each other except for the

common introduction (S, a) and end (f, E). The pupils do not begin at the same checkpoint – e.g. some start by walking to the lake, others to the forest.

Inspecting the graph reveals that the plots are made out of four self-sufficient episodes - the node pairs (b, 1), (c, 2), (d, 3), and (e, 4). The narrative arc of the play peaks for each node pair which makes the dramaturgical model an episodic wave model with four smaller climaxes. This narrative model makes it indifferent in which order the checkpoints are experienced, since their ordering depends on the path from node a. However, the narrative is written such that the order of the checkpoints and coordination nodes is fixed once the group is assigned a starting point (Figure 6).

At the end of the play, all of the pupils hear the same audio sequence. Though, the girls and boys get different versions of the same content of the audio files - that is, featuring a girl's voice for the girls, and a boy's voice for the boys. This is carried out in order to make it easier for the children to identify with the characters in the play. The theatre finds identification very important for the overall experience.

The incentive for combining dramatic courses and classes is to strengthen the learning possibilities by stimulating more senses; and to stage the learning in another and more exiting and motivating setting utilizing nature as the classroom.

Working with Mobile Urban Dramas for nature areas and not for cityscapes turned out not to be a limitation but rather an advantage. A forest or a lake can also be very dramatic settings; and when the children walk through the forest and hear the plaintively voices of the trees in the play, it is an advantage that the voices seem to emanate from one tree for one pupil, and from another tree for another pupil, depending on the position of the children.

4 Related Work

The three Mobile Urban Dramas developed with our framework represent a novel contribution to the domain of reality theatre/mobile drama. We briefly compare the Mobile Urban Dramas with state of the art. Earlier examples of urban storytelling include PodWalks [14] and AudioWalks [2], where the user downloads audio files to her mp3 player and listens to a narrative as she walks in the city. While these narratives refer to elements in the urban environment, there is no system support for integrating the city in the play's scenography and the stories are not interactive.

The "Call Cutta" [15] mobile phone play uses physical rooms in the city as scenography, and when the audience (a user) arrives, a phone rings and the play takes place as a partly improvised conversation between a real call centre person (actor) and the user. After a scene the user may continue to another room and have a new partly improvised conversation with another call centre actor. Compared to "Call Cutta" our concept provides a coherent storyline following a branching plot graph connecting a path through the city as the scenography. The user listens to the (branching) storyline while moving between the Semacode tagged locations that may involve some real actor activity, inspection of physical things, etc. in conjunction with a transition to a new audio scene in the headphones.

"Cargo Sofia" [4] has a storyline that develops relative to the route and locations visited, but the audience experience it only through a video projection shown in the dark container room, which resembles the traditional theatre arena, only smaller and

driving around in the streets. Even though there are similarities to Mobile Urban Dramas, “Cargo Sofia” does not utilize the native city locations as scenography, the city is experienced only through a video feed on the wall, and there are no possibilities for individual (e.g. branching) experiences or an individual pace on the route.

“Water” [18] is categorized as augmented reality fiction. Equipped with a notebook, GPS-System and earphones, the user walks to explore a story, which is determined by their physical movement. Like the mobile urban drama, “Water” aims to create a strong relation between the story and the surroundings. However, “Water” only explores single input (GPS) and output (audio from a notebook) and not the generalised concept that mobile urban drama incorporates.

Song Fuer C [19] is probably the project that comes closest to our Mobile Urban Dramas, but there is to our knowledge no usage of location-based technology, that allow individual branching or individual paces on the routes coupled to the storyline.

The New York eRuv [7] path provides a similar Semacode tagged path but with more factual information linked to the tags, and no branching story line. Thus it boils down to a location-based audio guide.

5 Conclusion and Future Work

We have developed a mobile phone software framework for realizing a novel dramaturgical genre – Mobile Urban Drama. The framework and the manuscripts developed by Katapult have given rise to a new genre of pervasive theatre experiences taking advantage of the user as the main character and utilizing the urban environment as the scenography. The framework is currently in use for several plays that have had several hundreds of paying theatre visitors. The experiences are generally very positive both from the audience and theatre reviewers from magazines and TV. We wish to develop the concepts further and we are taking advantage of the real use experiences in the development of new elements to the play support.

These development plans include a general reconnaissance of what the possibilities are – technical as well as dramaturgical – when designing mobile urban dramas. Further, the plans include a larger degree of user influence; e.g. giving the users the opportunity to make more elaborate and detailed choices in the story line. Moreover, different types of sensors and actuators are being integrated in the environment to activate the physical surroundings of the play - such as the buildings and interior of the city – to become interactive. In this manner, the surroundings will to a greater extent become an active part of the play, e.g. by means of slamming doors, showers of water, ambient sounds etc. This may give the audience interactive experiences using the city as an interactive theatrical scenography.

Acknowledgements. We thank our colleagues at www.interactivespaces.net. We would also like to thank the theatre Katapult, in particular Torben Dahl, Eva Mikkelsen and Solrun Toft Iversen, as well as the writers of the three plays, Michael Monberg, Bjørn Rasmussen and Mia Brandstrup Andersen, respectively. Further, we would like to thank for the support from VisitHorsens and HasleInteractive.

References

1. 4D-film (September-08-2008), http://en.wikipedia.org/wiki/4-D_film
2. AudioWalks (September-08-2008), <http://www.flamingoeffekten.dk/#41673>
3. Bernstein, M.: Storyspace 1. In: Blustein, J. (ed.) Proceedings of the Thirteenth ACM Conference on Hypertext and Hypermedia, HYPERTEXT 2002, College Park, Maryland, USA, June 11 - 15, 2002, pp. 172–181. ACM, New York (2002)
4. Brady, S.: Dicking Cargo Sofia – A Bulgarian Truck Ride through Dublin. In: TDR (The Drama Review), vol. 51(4) (T196), pp. 162–167 (Winter 2007)
5. Brusilovsky, P.: User Modeling and User-Adapted Interaction. Adaptive hypermedia, Ten Year Anniversary Issue 11, 87–110 (2001)
6. Cavazza, M., Pizzi, D.: Narratology for Interactive Storytelling: A Critical Introduction. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 72–83. Springer, Heidelberg (2006)
7. eRUV. A Street History in Semacode (September-08-2008), <http://www.dziga.com/eruv/index.php>
8. Hansen, F.A., Grønbaek, K.: Social web applications in the city: a lightweight infrastructure for urban computing. In: Proceedings of the Nineteenth ACM Conference on Hypertext and Hypermedia, Pittsburgh, PA, USA, June 19 - 21, 2008, pp. 175–180 (2008)
9. Hansen, F.A., Bouvin, N.O., Christensen, B.G., Grønbaek, K., Pedersen, T.B., Gagach, J.: Integrating the Web and the World: Contextual Trails on the Move. In: Proceedings of the 15th ACM Hypertext 2004 Conference, Santa Cruz, USA (2004)
10. Katapult Theatre (only in Danish) (September-08-2008), <http://www.katapult.dk/>
11. Kobsa, A.: User modelling: Recent work, prospects and hazards. In: Schneider-Hufschmidt, M., Kühme, T., Malinowski, U. (eds.) Adaptive User Interfaces: Principles and Practice, pp. 111–128. North-Holland, Amsterdam (1993)
12. Mateas, M., Stern, A.: Integrating plot, character and natural language processing in the interactive drama Façade. In: Gobel, S., et al. (eds.) Proceedings of the 1st Intl. Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2003), March 24-26. Fraunhofer IRB Verlag, Darmstadt (2003)
13. Mateas, M., Stern, A.: Façade, an experiment in building a fully-realized interactive drama. In: Game Developers Conference (GDC 2003), San Jose, CA, USA, March 4-8 (2003)
14. PodWalks (September-08-2008), <http://www.cphx.dk/index.php?id=34124>
15. Rimini-Protokoll (September 2008), http://www.rimini-protokoll.de/website/en/project_2766.html
16. Röber, N., Huber, C., Hartmann, K., Feustel, M., Masuch, M.: Interactive Audiobooks: Combining Narratives with Game Elements. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 358–369. Springer, Heidelberg (2006)
17. Ryan, M.-L.: Avatars of Story. University of Minnesota Press, Minneapolis (2006)
18. Schemat, S.: Water (Augmented Reality Fiction) (2004) (September-08-2008), http://www.kunstverein-cuxhaven.de/ohne_schnur/englisch/unterseiten/schemat.htm
19. Song Fuer, C.: Homepage (September-08-2008), <http://www.songfuerc.de/>
20. Szilas, N.: Stepping into the Interactive Drama. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 14–25. Springer, Heidelberg (2004)

Say Anything: A Massively Collaborative Open Domain Story Writing Companion

Reid Swanson and Andrew S. Gordon

The Institute for Creative Technologies, The University of Southern California
Marina Del Rey, Ca, 90292 USA
{swansonr, gordon}@ict.usc.edu

Abstract. Interactive storytelling is an interesting cross-disciplinary area that has importance in research as well as entertainment. In this paper we explore a new area of interactive storytelling that blurs the line between traditional interactive fiction and collaborative writing. We present a system where the user and computer take turns in writing sentences of a fictional narrative. Sentences contributed by the computer are selected from a collection of millions of stories extracted from Internet weblogs. By leveraging the large amounts of personal narrative content available on the web, we show that even with a simple approach our system can produce compelling stories with our users.

Keywords: Interactive Storytelling, Interactive Fiction, Collaborative Writing, Social Media, Information Retrieval.

1 Introduction

Interactive storytelling and interactive fiction on the computer have blossomed into active fields of research. Early work in these fields provided access to virtual worlds through textual descriptions, for example *Adventure* [1]. Although the level of sophistication varied, natural language was used to interact and learn about the underlying narrative structure. Today most state-of-the-art systems provide fully immersive 3D environments to depict what was once only abstractly available through textual responses [2][3][4][5]. Although the interface of modern systems is significantly different than their predecessors, many of the underlying goals are the same. Ultimately, one would like a system in which two competing and somewhat antithetical propositions are upheld. The human user should be free to explore and do anything in the world, receiving appropriate responses as they go. At the same time, a coherent underlying narrative should also be maintained in order to provide structure and meaning [6].

While much of the interactive media world has embraced advanced graphics technology for visualizing stunning virtual worlds, there are many reasons why purely textual works are still a valuable form of entertainment and a testbed for research. Although often scorned by both game developers and literary experts, Montfort [7] argues for the importance of textual interactive fiction on several levels. These include being a platform for computational linguistics research and improving language

skills and reading comprehension for the user. There also seems to be a void in the entertainment landscape between language games such as crossword puzzles and scrabble and graphic-rich video games.

Collaborative fiction is a particularly interesting literary exercise and topic to pursue because it is one of the few exceptions in traditional media in which the reader can also be an active participant in the shaping and unfolding of the narrative. Unlike traditional literature, which is perceived as a very personal act of creation with strong (single) authorship, collaborative works share the role of author amongst many individuals. Each individual contributes a portion of the story but the resulting piece has a life beyond any one of the creators. Despite the prominence of single-author assignments, the history of collaborative fiction is as long as ancient texts, such as the *Illiad* and the *Old Testament*. Collaborative fiction also has avenues for generating narrative works. For example, role playing games such as Dungeons & Dragons [8] are often seen as a process to generate narratives that arise through each character's decisions.

Computers and the Web have greatly facilitated communication amongst large numbers of people, which has led to an explosion of collaborative fiction between individuals. Not only has the Internet greatly assisted and expanded upon older collaborative writing genres, it has also spawned a variety of unique genres in its own right. Wiki-novels [9] and hypertext [10] are a few of the more popular examples. Rettberg [11] gives a more detailed introduction to these and other collaborative fiction genres.

The type of collaborative fiction we are interested in throughout this paper, however, is much closer to a traditional notion where people take turns writing segments of a story. This process can be an open-ended endeavor with no restrictions on content, length or structure but in many cases rules can be specified to constrain the work. Rules can pertain to the syntax of the story (such as the maximum number of words per turn or fixing the point of view of the story). Rules can also enforce semantic constraints like specifying a particular genre or requiring characters to have certain characteristics. Other types of rules can also be specified and are only limited by the imagination of the designers.

An important theoretic belief of this work is that when individual contributions are put together, a collective wisdom emerges that produces interesting relationships and properties that can transcend single authorship. With the introduction of the Web, vast amounts of people have easy access to each other with inexpensive means of publication. This has pushed the ideology of collective wisdom farther than ever before.

In this paper we will present a new type of interactive storytelling system that blurs the line between traditional interactive fiction and collaborative writing. We do this by leveraging the enormous amounts of content authored by ordinary people with a system that uses this content to take turns writing sentences with a single user. Potentially, millions of bloggers work together with the user of our system in a collaborative writing process to construct a new narrative work. We call our system *Say Anything* because, in the spirit of Bates' vision "Go anywhere and do anything" [2], the user is completely unconstrained in the topics about which they can write.

2 Say Anything

The primary virtue of our system is the simplicity in which our solution achieves its results. When the user contributes a sentence to the emerging story, we simply

identify the closest matching sentence in a large database of stories and return the following sentence of *that* story. With virtually no other preprocessing or modeling, relatively high-quality stories can emerge. For example, consider the following story segment created by a user of the *Say Anything* system:

You'll never believe what happened last night! Leigh laughed at my joke but I couldn't help but think 'liz would have laughed harder.' The joke wasn't very funny in a "ha ha" kind of way.. It wasn't anything like that, I thought he was going to give me a good night kiss but he ended up licking my cheek," she declared. It made me sneeze and snort out loud. And now my nose hurts from the snorting.

Instead of relying heavily on narratology and other narrative theories often used for interactive storytelling [12], we rely on emergent properties to give the story its structure, feeling and style. Although other genres, such as role-playing games, also justify their narrative structure on emergent properties, we believe our approach is somewhat more justified on this reliance. Many of the expected difficulties of controlling story structure have been eradicated because humans drive our story knowledge base and writing process. However, we concede that most authors on weblogs are not literary masters and the degree to which good narrative theory is available in our corpus is limited.

2.1 Story Corpus

To give *Say Anything* its generative power, a large corpus of stories is required. For this work we used the weblog story corpus developed by Gordon et al. [13]. The amount of user-generated content on the Web is rapidly growing and it is estimated that over 70 million weblogs exist [14]. Although written by ordinary people, not all weblog content is relevant to our goal of acquiring narrative content. Gordon et al. [13] estimated that only 17% of weblog writing is actually story content (descriptions and interpretations of causally related past events). The remaining text consists of political and other commentary, lists, opinions, quotations and other non-narrative subject matter. In their work they developed several automatic approaches for extracting the story text from these weblogs that incorporated bag-of-words, part-of-speech, kernel filtering methods and machine learning techniques. Since complete stories were more valuable than excluding non-story content, a version of their system that favored recall over precision was applied to a corpus of 3.4 million weblog entries. This identified 1.06 billion words of story content. Some post processing was necessary in order to sentence delimit the text and remove sentence fragments that were created as a result of the extraction process. The resulting story collection consisted of 3.7 million story segments for a total of 66.5 million sentences.

2.2 Sentence Retrieval

The fundamental mechanism of *Say Anything* is the identification and retrieval of similar stories in our database to the user's input. Contemporary information retrieval techniques provide exactly what we are looking for. In this work we used Apache Lucene [15], a freely available open source search engine. Lucene's algorithm

combines standard Boolean indexing with term frequency-inverse document frequency (TF-IDF) scoring functionality. An index is basically a lookup table whose keys are all the unique words/tokens contained in the document collection and the value is the list of all documents that contain that word. Various algorithms exist that allow you to query this index using Boolean operators efficiently. Given the set of documents these words appear in, TF-IDF is a common strategy for scoring and ranking them. All things being equal, words appearing frequently in a document should be given more significance. However, all things are not equal, and words that appear in many documents are given less weight because they are unlikely to be distinguishing words.

To implement this phase of the system, we merely need to treat each of the 66.5 million sentences as a document to be indexed. Simple tokenization is applied and each word is treated as a token in a bag-of-words approach. We then find the highest ranked sentence in our database and return the next sentence of that story.

2.3 The Interface

To complete the system, we needed a user interface. For this project we felt that a web-based interface would be most appropriate since the specifications were rather simple and no special graphics technology was needed. Going this route also allows good cross-platform compatibility and precludes the user from having to install special software on his or her machine.

The interface is split into two main regions. At the top is a panel for viewing the story, which is initially empty. At the bottom are controls for writing and appending sentences to the story. The process starts with the user writing a sentence to begin the story. Given what the user has typed, the system returns the next sentence in the emerging story. When the user feels the story has reached a good concluding point or the story has deviated too far off track (which is often the case), the user may end the story. At this point the user is asked to rate the story on two criteria using a 5 point scale:

1. *Coherence*: Do the system-generated sentences follow from what the user has written? Is the story coherent as a whole?
2. *Entertainment*: Did you have fun with the system? Was the story interesting or entertaining?

In order to successfully complete the story, the user must assign it a title.

3 Analysis

For our initial evaluation we set up a server on our lab's internal network to run our system. We recruited 11 people who wrote a total of 27 stories with the system over a period of two weeks. Our population consisted of men and women who worked in and around our lab, as well as a few other interested parties. We gave the users very little instructions except that they must write and rate their story with the system. We believed at this stage of our research that we should not limit the users to our preconceived notion of how the system should be used; moreover we hoped creative story writers could find functionality beyond our intentions.




Fig. 1. The number of ratings 11 users gave on 27 stories for the entertainment of the system (left) and the coherence of their story (right)

car chase

The user's sentence	Matching sentence	Returned sentence
On one sunny afternoon, I was driving down the highway.	cut to: so i was driving down the 110-S, one sunny L.A. december afternoon- i don't remember where i was going, but i was on the freeway in traffic.	and then i heard it.
An police car in my rear view mirror with its siren on full blast	Then he looked in his rear view mirror and saw a Police Patrol Car behind him, blue lights flashing and siren blaring.	I can get away from him with no problem thought the man and he tromped it some more and flew down the road at over 100 mph.
but his car was no match for the souped police interceptor.	He had his mother's car, a souped-up Ford Mustang.	At speeds topping 100 mph, he made for the airport.
But that proved out to be the wrong thing to do.	2. Thing that makes you feel great is: knowing that people that said i couldn't do it, are getting proved wrong 1.	I filled out 200 questions because: waitin on fatty to get here so i can leave!
yes I was at the police station witing for my lawyer to bail me out.	I thought it was hilarious when David showed up at the police station with the money to bail me out of jail.	Ha!
It is never a good idea to run from the law, atleast not in a 100hp car.	Last run in with the Law: never.	knock on wood 11.

Fig. 2. An example of a story, titled *car chase*, created with the system and given a coherence rating of 4. On the left are the sentences the user submitted to the system. In the middle are the highest ranked sentences in the database and on the right are the next sentences returned to the user by the system.

There was a wide variety in story generation from the various users. The shortest story consisted of only 2 sentences (1 turn), the longest 30 sentences and an average of 12.7. On average the users had a positive experience using the system 3.39 (1.23 stdev) on a rating scale from 1 (low) to 5 (high). Generally the users did not believe the stories were very coherent. The average coherence rating was 2.29 (1.24 stdev) also on a 5 point scale. See Fig. 1 for a more complete summary of the ratings.

3.1 When Things Go Well

Although the ratings give some indication about the quality of the system, the small sample size and inability to get a feel for what was written limited our knowledge about the system's capabilities. It is more informative to examine some sample output. Fig. 2 presents a complete story from the collection. In some cases it is striking how close the user's sentence matches one from the database (for example those in the first two turns). Despite the poor matching in the third turn, the story does not lose continuity until the fourth turn. Even here, however, the user is able to make the story fairly coherent by shifting the setting to the police station. Although there are some minor peculiarities, the story remains on track until the end. Without any explicit model underneath to guide the process, this story is able to produce a fairly regular narrative structure—an introduction, scene setting, plot development and a climax. The final narrative element, denouement, occurs with the character realizing the errors of his ways. This example also shows that the process is fairly robust to spelling errors.

3.2 When Things Do Not Go As Well

The previous example gives a feel for the kinds of stories that can be produced and how the system works, however, there are several issues that have not been highlighted. An excerpt from a different story, illustrated in Fig. 3, reveals several of the prominent problems with the system.

As is the case for most of the stories, the two major problems are coreference resolution and event ordering/prediction. It is rather impressive that the coreference between entities in *car chase* can be interpreted correctly throughout the story. This is remarkable because at this stage of development no special care has been devoted to them. However, as seen by the example in Fig. 3, this is not always the case. The first sentence in this sequence sets the stage for a driving event that takes place in summer. For coherence, the verbs subscripted with the number 1 would make most sense if they were all part of the main driving event, or related to it in some way. However, in the matched sentence the action *drinking* (wine) is completely unrelated. This is worrisome because it increases the likelihood that the following sentence from that story will be unrelated to events that often occur while driving. Fortunately the next sentence, the one returned to the user, is generic enough that it does not completely spoil the coherency. It is not until the character is sledding that repairing the story becomes nearly impossible. Similarly the phrases subscripted with the number 2 should all refer to the same object, namely the *happy pills*. However, since the retrieval strategy only considers one sentence, there is no way during this simple query to unify the objects. In this case it causes the pills to change into mittens and a sled. This also seems to be the reason that the main event, driving, shifts to become a *sledding* event.

This example also illustrates how difficult natural language can be when dealing with implied or implicit knowledge. This can be seen with the explicit reference to the phrase *middle of summer* as indicated by the subscript number 3. Although it is not explicitly indicated, there is an implied reference to winter later on when the main character is sledding. This is indicated in the illustration by an implicit phrase [*in winter*], which does not appear in the actual story. A creative writer could probably work his or her way around this change, although two major shifts in expectation in such a short duration are a lot to overcome.

In addition to these issues, several other problems also frequently arise. Similar to the coreference problem is that of number, gender and case agreement. Some problems also arise specifically out of the nature of the ranking algorithm. Since TF-IDF only considers the frequency of terms, some unfortunate side effects occur. For example, without any preprocessing, numbers and other rare tokens pose a significant problem as seen by the input “*She lived in apartment 4311. I spend many pleasant nights there.*” entered by the user¹. There are many things one would hope the system would pick up on to choose the next sentence. However, what the TF-IDF algorithm deemed the most similar was the number 4311. This is most likely due to the fact only one sentence in the database contained this number giving it a very high inverse document frequency score. However, this token is only rare by coincidence and not because it is meaningful in any way. Modifying words and phrases also significantly impacts the ability for the retrieval to stay on track. For example, negation, statements of belief or qualifying statements have significant semantic importance but often do not play an important role in finding an appropriate matching sentence.

Although some of the problems seen above occur directly from the IR techniques at the sentence level, the much bigger issue is how to ensure that the global features of the story remain consistent throughout. As demonstrated above, the most difficult issues appear to be event sequence compatibility and ensuring named entities refer to the things they are supposed to. This is not an easy problem to solve but solutions may include event prediction and modeling [16][17][18] and integrating coreference resolution techniques [19].

The user's sentence	Matching sentence	Returned sentence
Anyway I should know better than to <i>drive</i> ₁ through red flag areas in the <i>middle of summer</i> ₃ .	I know better than to <i>drink</i> ₁ red wine after ten, but I did it anyway.	I found myself wide away at four AM with little pain in my face, so I took two of my <i>happy feel good pills</i> ₂ .
I wished I had <i>some</i> ₂ right then.	I wish I had a <i>sled and some mittens</i> ₂ .	I would have been out there with the kids <i>sledding</i> ₁ anyway [<i>in the winter</i>] ₃ .

Fig. 3. Two turns from a story that show problems that can easily lead stories of their path. Note that [*in winter*] does not actually appear in the user's story.

¹ This user broke the rules slightly by inputting more than one sentence at a time.

4 Discussion

Achieving an open domain system that is capable of producing high-quality stories is an extremely difficult undertaking. One of the common criticisms of traditional interactive storytelling mechanisms is the limited domain in which the user can interact with them. Even when the domain is heavily restricted, the cost associated with authoring the content is considerably expensive. For example, Façade [20] one of the premier systems today took two man years to author the content. Although the resulting content and system are among the best, the simulation is only playable for 20 minutes 6 or 7 times. The breadth of coverage and rapid development time are highly attractive despite there being a long way to go before our system produces significant, well structured stories.

Although similar in many respects, our system differs from prototypical interactive storytelling systems in a few key ways. Our goal is not to guide agents in a virtual world in order to tell a pre-authored story, but instead to provide a virtual collaborative writing environment in which the human and computer take turns authoring an entirely new story. Also most interactive story systems take a top down approach where high-level knowledge is hand encoded. In our system we take more of a bottom-up approach where knowledge is acquired from a large corpus of existing stories. This approach allows us to scale to any domain covered in our database and removes many of the difficult narrative theory concerns from the system architecture. In a traditional approach where individual experts write the content for each domain, scaling to the breadth of topics that are common in every day life will be prohibitively expensive. Instead, our system leverages the massive amounts of users who provide this knowledge, not in any logical form, but in plain English that takes no special skill to contribute.

5 Conclusions

In this paper we have presented a new type of interactive storytelling architecture that is unique both in its user interaction model and in the way the system models the story generation process. The sample output and coherence ratings show that there is still a long way to go before high quality stories are capable of being produced. However, there are several encouraging factors that underscore the promise of this approach. One of the most attractive aspects of this architecture is the time of content development. On the one hand the number of man-years required to author all of the content is enormous, consisting of tens-of-millions of weblog users who continue to provide more data than can be reasonably integrated into this system. Although the stories meander without a strong direction more often than not, this methodology shows it is possible that a decent narrative structure can be accomplished using a data-driven model, especially with continued refinement.

Acknowledgements. The project or effort described here has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

1. Crowther, W., Woods, D.: *Adventure* (1977)
2. Bates, J.: *The nature of characters in interactive worlds and the oz project*. Technical report, Carnegie Mellon University, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA (October 1992)
3. Cavazza, M., Lugin, J.L., Pizzi, D., Charles, F.: *Madame bovary on the holodeck: immersive interactive storytelling*. In: *MULTIMEDIA 2007: Proceedings of the 15th international conference on Multimedia*, pp. 651–660. ACM, New York (2007)
4. Riedl, M.O., Young, R.M.: *An objective character believability evaluation procedure for multi-agent story generation systems*, pp. 278–291 (2005)
5. Swartout, W., Gratch, J., Hill, R.W., Hovy, E., Marsella, S., Rickel, J., Traum, D.: *Toward virtual humans*. *AI Mag.* 27(2), 96–108 (2006)
6. Bates, J.: *Virtual reality, art, and entertainment*. *Presence: Teleoper. Virtual Environ* 1(1), 133–138 (1992)
7. Montfort, N.: *Generating narrative variation in interactive fiction*. PhD thesis, Philadelphia, PA, USA (Adviser-Mitchell P. Marcus and Adviser-Gerald Prince) (2007)
8. *Dungeons & Dragons official website* (2008), <http://www.wizards.com/>
9. *Wikinovels* (2008), <http://bluwiki.com/go/Wikinovels>
10. Konklin, J.: *Hypertext: An introduction and survey*. *Computer* 20(9), 17–41 (1987)
11. Rettberg, S.: *All together now: Collective knowledge, collective narratives, and architectures of participation*. In: *Digital Arts and Culture* (2005)
12. Cavazza, M., Pizzi, D.: *Narratology for interactive storytelling: A critical introduction*. In: *Technologies for Interactive Digital Storytelling and Entertainment*, pp. 72–83 (2006)
13. Gordon, A.S., Cao, Q., Swanson, R.: *Automated story capture from internet weblogs*. In: *K-CAP 2007: Proceedings of the 4th international conference on Knowledge capture*, pp. 167–168. ACM, New York (2007)
14. *Technorati State of the Blogosphere / State of the Live Web* (2007), <http://www.sifry.com/stateoftheliveweb>
15. Hatcher, E., Gospodnetic, O.: *Lucene in action*. Paperback (2004)
16. Bejan, C.A.: *Unsupervised discovery of event scenarios from texts*. In: *The Florida Artificial Intelligence Research Society*, Coconut Grove, FL (May 2008)
17. Manshadi, M., Swanson, R., Gordon, A.S.: *Learning a probabilistic model of event sequences from internet weblog stories*. In: *The Florida Artificial Intelligence Research Society*, Coconut Grove, FL (May 2008)
18. Chambers, N., Jurafsky, D.: *Unsupervised learning of narrative event chains*. In: *Proceedings of ACL 2008: HLT*, pp. 789–797. Association for Computational Linguistics, Columbus (2008)
19. Elsner, M., Charniak, E.: *Coreference-inspired coherence modeling*. In: *Proceedings of ACL 2008: HLT, Short Papers*, pp. 41–44. Association for Computational Linguistics, Columbus (2008)
20. Mateas, M., Stern, A.: *Façade: An experiment in building a fully-realized interactive drama*. In: *Game Developer's Conference*, San Jose, California (March 2003)

Locating Drama: A Demonstration of Location-Aware Audio Drama

Nye Parry, Helen Bendon, Stephen Boyd Davis, and Magnus Moar

Lansdown Centre for Electronic Arts
Middlesex University, Cat Hill, Barnet, Herts EN4 8HT
United Kingdom

nye@nyeparry.com, {h.bendon,s.boyd-davis,m.moar}@mdx.ac.uk

Abstract. Locating Drama is a collaborative project between the Lansdown Centre for Electronic Arts at Middlesex University and the BBC Radio Drama department. The aim of the project is to investigate narrative strategies that utilise locative technologies, principally GPS enabled devices, without being tied to a specific location, allowing listeners to experience immersive, location-aware (but non-location-specific) dramas in suitable locations near their homes, which may eventually be downloaded as interactive pod-casts from the BBC website. The demonstration will present a working drama originally produced for the BBC's Free Thinking festival 2008.

Keywords: GPS, Locative Media, Audio Drama, Interactivity, Narrative.

1 The Project

Locating Drama is a locative audio drama project developed jointly by members of the Lansdown Centre for Electronic Arts at Middlesex University and the BBC Radio Drama department. Locative media offer an important new field in which the physical location of listeners affects their listening experience – location-aware media devices are rapidly becoming available in the form of GPS-enabled PDAs and mobile phones.

The aim of the project is to investigate narrative strategies that utilise these technologies without being tied to a specific location, allowing listeners to experience immersive, location-aware (but non-location-specific) dramas in suitable locations near their homes, which may eventually be downloaded as interactive pod-casts from the BBC website.

The current phase of the project is focused on trials to be carried out at the BBC Radio 3 Free Thinking festival in Liverpool in early October 2008 for which a working interactive drama will be produced. At the time of writing preliminary trials have been carried out to test different “story-spaces”, transitions and narrative voices. A working drama production will now be produced for Free Thinking 08 which we will also be able to demonstrate practically at the Interactive Storytelling conference.

As the focus of the research is on scalability and transportability of locative media experiences we will present a complete working version of the drama to the conference, to demonstrate on GPS enabled PDAs.

2 Context

As Iain Chambers observes, the advent of the personal stereo has changed not only where and when we listen, but also how we relate to what we hear: “Here as opposed to the discarded “grand narratives” (Lyotard) of the city, the Walkman offers the possibility of a micro-narrative, a customised story and soundtrack, not merely a space but a place, a site of dwelling. Our listening acts as an escape from our lived environment while also intersecting with this environment forming an accidental soundtrack to our real lives.” [1] This leads to the possibility of what Siegfried calls “situational stories” [2], in which both the listener and the listening environment itself can be considered active participants in the story-telling.

While the majority of projects that use GPS in a narrative context (Moving Audio’s 230 Miles of Love or the BBC’s Riot! 1831 and Culloden) have focused on the exploration of a particular site and some (such as those by Blast Theory) have even incorporated live actors into the narrative, making them both site- and time-specific, our project aims to open locative media experiences to a wider public through a location-aware approach that is not tied to a specific map.

GPS supplies information about the absolute position of the user as well as the date and time. It may also, by implication, be used to respond to relative position: how far has a user moved in a given time, has the user returned to their starting location etc. It is this information that we use as a starting point for interaction. It is possible to calculate approximate speed of movement, direction of movement, shape and extent of a journey, for example, and use this information to generate a user-specific mix of previously authored content. This allows a number of fascinating narrative possibilities: characters that you meet in a particular location each time you visit it, voices calling out, demanding attention which may be selected by stopping to listen or rejected by moving on. A significant feature of this type of interaction is that the user is not obliged to interrupt the narrative experience by looking at the device, interacting using buttons and stylus etc. Instead changes in the experience are prompted simply by moving about. The interaction in that sense is ‘transparent’ to the experience. Our previous experience confirms that this is an important benefit for users.

Another concern of the project was highlighted at a preliminary open discussion at Free Thinking 07, which launched the project. It emerged that most BBC staff identified locative media with games rather than serious drama, and in several cases with user-generated content. This project instead is concerned with the relationship between locative media and the writing and production of drama by experienced professionals. A listener-responsive experience, it is not about user-generated content but a form of storytelling, offering audiences the opportunity of interacting with fragmented narratives within any physical location. The potential for access by a mass audience is important.

3 Trials

We have undertaken a series of content trials on GPS devices to explore appropriate narrative forms for the medium and user engagement. These trials explore both technical possibilities and narrative structure, to begin to get a sense of the possibilities for the direction of the project.

A number of conclusions can already be drawn from these early trials.

1. It is important to clearly define a “story-space” in which the action takes place. Trials which used defined boundaries between virtual spaces were particularly successful. In one case a ring of applause defined the border of the story-space on entry (replaced boos on exit). Similarly a clear boundary (sonic curtain) between a contemporary and a historical space made an effective transition between alternative narrative timelines.
2. Silence is difficult. Allowing users to drop out of the experience between different zones of action is disconcerting and causes some confusion. A continuous sound design strategy has been adopted within the story-space utilising cross-fades as well as more abrupt transitions to set different scenes.
3. Revisiting locations and characters allows a sophisticated picture of events to be built up. Characters can be associated with a place which once visited can be returned to for more information. This may be used to create a voice of authority, a home or safe place, though other characters may also cast doubt over the legitimacy of this voice.
4. Instructions should, where possible, be embedded in the narrative. The experience should be seamless and avoid the need for looking at a screen or using any interface device beyond the listener’s movement. Specific instructions may take the form of “voices in your head” or external characters calling out or advising you to return to a location you have previously visited.

These ideas will inform the writing and production process ahead of the Liverpool presentation which will be undertaken collaboratively between the BBC and Lansdown Centre teams. The project is being implemented using Hewlett-Packard’s Mediascape software and Flash and will run under Windows on any GPS-enabled PDA. Thanks are due to the BBC and to Hewlett-Packard for the supply of demonstration equipment.

References

1. Chambers, I.: The Aural Walk. In: Cox, C., Warner, D. (eds.) *Audio Culture, Readings in Modern Music*, Continuum, New York, p. 100 (2004)
2. Siegfried, W.: *Soundtracks to Reality* (1996), <http://www.ariarium.de/stracks.htm>

Lies and Seductions

Petri Lankoski¹ and Tommi Horttana²

¹University of Art and Design Helsinki, Hämeentie 135C, FI-00560
Helsinki, Finland

petri.lankoski@taik.fi

²Helsinki University of Technology, PL 1000, 02015 TKK, Finland

tommi.horttana@tkk.fi

Abstract. *Lies and Seductions* is a computer game in which a player controls Abby, a character on a wager to seduce a rock star who has promised to stay a virgin until marriage. The game is loosely based on the story of *Les Liaisons dangereuses* (*Dangerous Liaisons*). The game is a design experiment exploring how to use social relations, seduction, and tragedy as meaningful gameplay content.

Keywords: Computer game, social interaction, seduction.

1 Introduction

Lies and Seductions^[5] is a 3D computer game about seduction, lies, betrayal, and tragedy. The game is built using Unity3d^[6] game development tool. The gameplay focuses on social interaction with non-player characters (NPCs). Information and dirty little secrets are gathered through discussion, gossip and eavesdropping to be able to navigate a social network to the heart of a rock star. The game draws from the novel *Les Liaisons dangereuses* (*Dangerous Liaisons*) by Pierre Choderlos de Laclos (first published in 1782) and its film adaptations *Dangerous Liaisons*, *Cruel Intentions*, and *Untold Scandal*.

With some liberties, we took the main characters and their relations, as a starting point for our game. The main challenge in the game design has been to turn the characters, their relations, and, especially, the main conflict into gameplay.

The main conflict we used in the *Lies and Seduction* is adapted from *Dangerous Liaisons*. Marquise de Merteuil and Vicomte de Valmont are decadent nobles and ex-lovers. De Valmont is a skilled womanizer who has set his eyes on the beautiful, married, and chaste Madame de Tourvel. The marquise makes a promise to have sex with de Valmont if he succeeds to seduce de Tourvel. The story of *Dangerous Liaisons* contains another seduction assignment for de Valmont, but we did not use it in the game.

Lies and Seductions game is a design experiment on character-driven game design. Gameplay design patterns “are semiformal interdependent descriptions of commonly reoccurring parts of the design of a game that concerns gameplay.”^[1] In the design of NPCs, the social networks and conflicts we have used are found in several patterns described by Lankoski and Björk.^{[3], [4]}

2 Characters and Gameplay in Lies and Seductions

The player character (PC) in *Lies and Seductions* is Abby (the character's counterpart in the *Dangerous Liaisons* is de Valmont). Abby and Becca (de Merteuil) are on a cruise. They make a bet: if before the cruise ends, Abby gets to seduce Chris (de Tourvel), the singer of a rock band, Becca sleeps with Abby. Otherwise, Abby lets Becca humiliate her publicly. The challenge is that Chris has publicly promised to stay virgin until marriage and has, so far, kept his word. We use a cut-scene to introduce the bet. Chris is traveling with his two friends, Emma and Ed. Ed is the bass player and songwriter in Chris's band. Another noteworthy character is rich aristocrat Lord James who wants to have sex with Abby. These characters serve as a starting point of the gameplay design.

The main goal in the game is to seduce Chris before the deadline. To complicate things, this is not possible directly, as Chris initially refuses to engage in longer conversations with Abby. A Player needs to convince one of Chris's friends to help. In this way there are two goals that regulate advancement (on regulating goals, c.f., Lankoski ^[2]); one of these goals needs to be accomplished before one can successfully reach the main goal of the game.




Fig. 1. Dance minigame

The above goal structure, as such, would create a rather static game structure. To make things more dynamic, we use a gameplay design patterned after *Actions Have Social Consequences* (c.f., Lankoski and Björk ^[4]). This means that Abby's actions will shape how a non-player character feels about Abby. For example, if the player guides Abby to kiss some other character, and Chris witnesses this, Chris will have a negative impression on Abby. This, consequently, will make seducing Chris more difficult. The player's performance in the minigame, for example in a dance mini game (Fig. 1), can also change the impressions of NPCs.

Design-wise, this is implemented using an impression system, in which certain actions set an impression in an NPC’s memory. In the above example, the impression “Slutty” is adjusted in Chris’s memory. This, along with other impressions, affects Chris’s overall attitude to Abby (meaning, in the context of the game, his willingness to have sex with her). How each impression affects each NPC’s attitude depends on their personality. For example, while Chris dislikes sluttiness, some other characters find that enticing. Using the impression system, it is also possible to add some level of *Memory of Important Events* (c.f., Lankoski and Björk ^[4]) for NPCs. Impressions can also be used to trigger *Emotional Attachment* (c.f., Lankoski and Björk ^[3]), which means that an NPC has and shows emotional reactions to certain events in the game.

Storing the impressions separately, instead of simply adjusting a single attitude stat for each NPC, allows the NPCs to act on individual impressions (for example, mention them in dialogue). More interestingly perhaps, it also allows them to share their impressions about Abby through gossip. This gossiping is *Information Passing* (c.f., Lankoski and Björk ^[4]) between NPCs, and this gossiping influences gameplay.



Fig. 2. A gameplay screen with dialogue

Conversations with NPCs have an important role in the game. We use prewritten dialogue trees, in which the availability of branches are based on the current impressions of an NPC. Impressions are used to create *Contextualized Conversational Responses* (c.f., Lankoski and Björk ^[3]). In addition, choices made by players during a conversation can change impressions. The discussion with an NPC is shown in figure 2 above. For writing the dialogue, we have built an editor that can also be used to simulate dialogues with impressions (see figure 3).

In conclusion, *Lies and Seductions* is a game in which we explore ways to bring some structures of drama, tragedy, and social interaction to the game. Design-wise, we have been interested in testing various ways of integrating scriptwriting and gameplay design seamlessly.




Fig. 3. Our dialogue editor and simulator

Acknowledgements. *Lies and Seductions* is a game developed by Petri Lankoski, Tommi Horttana, Niklas Gustaffson, Kimmo Karvinen, Reeta-kaisa Neittaanmäki, Linda Kronman, Anne Parkkali, Raisa Omaheimo, Mike Pohjola, and Jari Suominen.

References

1. Björk, S., Holopainen, J.: *Patterns in Game Design*. Charles River Media, Hingham (2005)
2. Lankoski, P.: *Goals, Affects, and Empathy in Games*. In: *Philosophy of Computer Games*. Reggio Emilia, Italy (2007)
3. Lankoski, P., Björk, S.: *Gameplay Design Patterns for Believable Non-Player Characters*. In: *Situated Play*. Digra Japan, Tokyo (2007)
4. Lankoski, P., Björk, S.: *Gameplay Design Patterns for Social Networks and Conflicts*. In: *GDTW 2007*. John Moores University, Liverpool (2007)
5. Lankoski, P., Horttana, T., Gustaffson, N., et al.: *Lies and Seductions* (2008)
6. Unity Technologies: *UNITY: Game Development Tool*, <http://www.unity3d.com>

Animation-Based Interactive Storytelling System

Kaoru Sumi

National Institute of Information and Communications Technology,
3-5 Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0289, Japan
Kaoru@nict.go.jp

Abstract. This paper introduces an animation story generation system called WordsAnime using the animation database provided by animation Consumer Generated Media. A user generates a story by repeating the steps of inputting scenarios and viewing animations. The system presents manual, semiautomatic, and fully automatic modes of creating scenarios using a common knowledge base that links different scenarios with each other.

1 Introduction

WordsAnime is a story generation system by using animation data collected by Anime de Blog [1]. Even children could create animation content easily by developing the story with WordsAnime. The system leads users to input the story line through the method of paragraph writing. If a user inputs a scenario, an animation corresponding to the scenario appears on the display. WordsAnime is software supporting story generation so that child users can have great time creating animation content.

Our previous system, Interactive e-Hon [2] is a media translation application that translates text into animation by linking animation data to words. Its open problem was that it requires preparing an enormous amount of animation data to match all the words that people are likely to use.

There are some technologies creating new content through content combination such as scripting languages [3][4][5][6][7][8]. Considering to offer greater flexibility for editing, they also have the problem of how to prepare enough animation data would remain.

Anime de Blog is a Consumer Generated Media that aims to collect animation data matching all the words that people usually use and WordsAnime is its application. Using Animebase of Anime de blog, we can use unlimited animation data corresponding to words. There are not many systems that enable users to create stories showing animation content and have interfaces with simple selection operations and language that can be used to describe a fair amount of animation data.

Works created by WordsAnime are stored on the WordsAnime server. WordsAnime displays GIF animation data for displaying on the screen from Animebase, stored on the Anime de Blog server. WordsAnime includes functions for story generation support, automatic scenario generation using rules, and story browsing. It is based on Servlet/JSP and Ajax (JavaScript) and requires Windows XP and Internet Explorer 6.0 SP2 or Firefox 2.0.0.9.



Fig. 1. Snapshot of “Introduction”

2 Animation-Based Interactive Storytelling System

The user generates the complete story by repeating the steps of inputting scenarios and viewing animations. WordsAnime uses a method of paragraph writing enabling users to easily create stories. The system leads a user to create scenes for “Introduction”(Figure 1), “Problems”, “Challenges”, “Problem Resolution” and “Conclusion”. The sequences from “Problems” to “Problem Resolution” and from “Problems” to “Challenges” are repeatable. For each scene, the system leads users to input “Who?” and “Doing what?” as necessary input items. All the input items can be selected by clicking alternatives in the scrolling list or typing the user’s own words. As unnecessary items, the system provides “When?”, “Where?”, “What?”, “Whom?”

WordsAnime has rules of a common knowledge base for automatic scenario generation. A user can describe rules by writing IF and THEN sentences. An IF sentence represents the condition of the scene. A THEN sentence represents the condition of the next scene. When the user creates a scenario and clicks the button “Finish the scene”, then some scenario candidates appear as hints for the user. The user can select one of the candidates or create his or her own scenario when creating a scene. If the user selects one scenario, then the corresponding animation appears.

Users can apply one of two modes for automatic scenario creation when they want to use automatic scenario creation. The first mode requires selecting one of the candidates in any one scene. The second mode is a fully automatic mode, in which the complete story plays without input from the next scenario until the end.

3 Discussion

WordsAnime enables content generation and transmission by novice users; even children can easily create stories on the Web. The mental images that people have when they create something vary from person to person. In our system, users can select parts of animation content from Animebase and combine them. Although they cannot always create content that precisely matches their own imaginations, they can come close to it by using our approach and large animation databases.

WordsAnime was developed as an application of Animebase, which is the animation database for Anime de Blog. Animebase was applied to many types of applications. Animebase will enable novices to create content by offering an environment in which the data can be freely shared. It will be useful for the following: education, e-learning based on digital content, understanding and learning assistance for children or the elderly, business, transmission of information that appeals to individuals with related services, advertising and marketing using digital content, communication support support for people of different regions, cultures or background knowledge.

We believe that new research and business opportunities will be created and new areas of information science will be pioneered by allowing the academic and business access to the database.

References

1. Sumi, K.: Anime Blog for collecting Animation Data. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 139–149. Springer, Heidelberg (2007)
2. Sumi, K., Tanaka, K.: Automatic Conversion from E-content into Virtual Storytelling. In: Subsol, G. (ed.) ICVS-VirtStory 2005. LNCS, vol. 3805, pp. 260–269. Springer, Heidelberg (2005)
3. Hayashi, M.: TVML(TV program Making Language) Make Your Own TV Programs on a PC! In: International Conferences, Virtual Studios And Virtual Production (2000)
4. Huang, Z., Eliens, A., Visser, C.: STEP: a Scripting Language for Embodied Agents. In: Prendinger, H., Ishizuka, M. (eds.) Life-Like Characters: Tools, Affective Functions, and Applications, pp. 87–110. Springer, Heidelberg (2004)
5. Cassell, J., Vilhjalmsson, H.H., Bickmore, T.: BEAT: The Behavior Expression Animation Toolkit. In: Prendinger, H., Ishizuka, M. (eds.) Life-Like Characters: Tools, Affective Functions, and Applications, pp. 163–186. Springer, Heidelberg (2004)
6. Prendinger, H., Saeyor, S., Ishizuka, M.: MPML and SCREAM: Scripting the Bodies and Minds of Life-Like Characters. In: Prendinger, H., Ishizuka, M. (eds.) Life-Like Characters: Tools, Affective Functions, and Applications, pp. 213–244. Springer, Heidelberg (2004)
7. Mariott, A., Stallo, J.: VHML-Uncertainties and problems. A discussion. In: Proceedings AAMAS 2002 Workshop on Embodied Conversational Agents - Let's specify and Evaluate them! (2002)
8. Microsoft: Developing for Microsoft Agent. Microsoft Press, Redmond (1998)

Dear Esther: An Interactive Ghost Story Built Using the Source Engine

Dan Pinchbeck

School of Creative Technologies, University of Portsmouth, UK
Dan.Pinchbeck@port.ac.uk

Abstract. This paper reflects on the design and production of an multimodal, environmental storytelling experiment constructed in the first-person game engine Source. Rather than being based around the resolution of conflicts and achieving goals, Dear Esther presents a sparse environment with no embedded agents, relying purely on the player's engagement with and interpretation of a narrative delivered through semi-randomised audio fragments. Dear Esther was released for free download via a number of modding sites in June and this paper reflects on the experience of building and the response by gamers.

1 Introduction

In 2007, we were awarded a speculative research grant from the Arts & Humanities Research Council, UK, to develop three mods, each exploring a different angle on storytelling, or affective structures, and release them freely on the internet, tracking their responses. We were not interested in creating mods that only had value as objects of study, but wanted our results to succeed as games in the public domain. In other words, rather than produce media which was primarily of interest to an artistic or academic audience, we wanted to release games that pushed the envelope whilst remaining appealing for gamers.

Dear Esther runs as a first-person game, with no visual representation of the avatar. There are no goals or action sequences, just an environment to explore with embedded music and voice-over triggers. Additionally, the game contains no AI, making the player's engagement with the piece rest entirely with the narrative, visual environment and audio. The environment is custom-built, but it uses texture files pre-existing in Source, plus a small number of new visual elements as required by the narrative. A new entity was coded which allowed us to fire one of three audio files for each trigger point. In this way, the story is effectively randomised. Dear Esther is superficially familiar as a ghost story, though it becomes more abstract and psychological as it progresses. Two plots develop simultaneously: the avatar's visit to a deserted island, following the text of a 17th century cartographer, and memories of a car accident. We begin to understand that whoever Esther is, she has been killed in the accident and the avatar has come to the island because he believes he is being drawn to find her. This is complicated by the introduction of Paul, the man he believes caused

the crash, although it is inferred that these memories also refer to the biblical character's conversion on the road to Damascus. This is indicative of the collapse of symbols within the story: his memories of having kidney stones leads to the inference that the island is itself some kind of metaphoric stone; the identity of a shepherd who died hundreds of years before becomes confused with the narrators; his job as an electrical engineer becomes an obsession with interconnections; the aerial that towers above the landscape becomes a conduit for redemption. Esthers identity is never revealed and ultimately, there remains the question of whether the island is real at all; whether the accident ever happened; whether the narrator himself exists in any real sense. Compounding this, the randomisation of the narrative fragments and contradictions coded into the text means a closed reading of the events is impossible to ever reach.

2 Development

The plot was developed with a consideration of exploratory diversity and the economics of building. It was important to have a structure that would encourage exploration and a drive to complete the experience. Conflict was discarded, along with agents. Taking the basic idea of a ghost story, which enabled a mysterious, empty landscape, with aspects of a psychological drama, where the player-avatar relationship would be pushed to the fore, we decided on a remote island as a backdrop to the story. Localisation, the separation of the game environment from the wider, inferred world, is a common strategy in games, as it allows a greater control to be taken of the contents of the presented environment. Researching islands to base the environment on also led us to historical documentation of Boreray, a Hebridean island. There were a number of advantages of using an island in this location: we could constrain exploration with believable geographical features (cliff edges, steep valleys and the ocean), use vertical features as focal points, and bring in a limited set of human markers (Boreray was a grazing point for lone shepherds). It also enabled us to make use of Source's fog effects to limit draw distances and cover-up the limitations of the environment's size, and the textures pre-existing in the engine. Finally, it allowed us control over visual diversity, so features would stand-out, increasing our ability to use visual focalisation. We established a focus point in an aerial on a mountain, giving the player an indication of their ultimate exploratory goal.

Freed of the need for microgoals and ludic set-peices, we had the opportunity to write and build simultaneously. Rough timing of exploration was used as a way of determining the number of plot triggers that could be used, at this point nearly fifty cues were embedded. At this stage, the overall plot arc had been determined, so it was a case of trying to ensure that triggers were evenly distributed, whilst keeping momentum. Visual symbols, picking up images in the textual fragments, were also built, ranging from cave paintings to creating parts of the geography to fit the narrative. What was telling from the first version was that the story was not simply competing for space with gameplay, but the combination of a rich environment with open exploration meant that the project

felt ‘overstoriaed’. Although trigger positions had been carefully worked out, the narrative fragments were more or less constant when running the direct path of exploration, resulting in a very uneven experience if the player chose to return to a prior space. Further to this, randomisation meant that the plot had become too abstract, shifting from any kind of recognisable narrative to a barrage of textual symbols. Thus, the number of plot points was reduced, and single triggers replaced random triggers in key locations. We also wanted to avoid versions of the same text at randomised triggers, and this meant that we to make sure that key points were not lost. Although the use of bottleneck triggers was kept as small as possible, it nevertheless represented a compromise and illustrated the limitation of completely free association with a coherent narrative line. Whilst we wanted to enable free exploration, we had to balance the triggers against one another to ensure we were not getting crossovers - firing one cue whilst another was already playing. Adjusting the speed the avatar moved at was one way of doing this, but had to be constrained by the playability of the experience: beyond a point, the whole thing became unnaturally sluggish. In a few instances, the environment was linearised, to force the player around a lengthier route; this created an environmental sub-plot structure, where sections of exploration were avoidable. A riskier strategy was to try and direct the player by using more explicit focalisation devices. On three occasions, a gull takes off near the player and flies in the direction they are being prompted to take. We also, very sparingly, used a human agent, always viewed at a distance, to serve a similar purpose. The importance of music also became more apparent: it was critical to establish a continuity of mood, to allow us to be more fragmented with the actual text. By embedding music cues with a common stylistic base, spanning multiple trigger points, a sense of linear development could underpin the interactive explorations. Additionally, music also helped to create an affective score which enabled us to manipulate the player’s emotional experience, orientation and position in the wider story arc, and even adjust exploration.

3 Response from the Community

Dear Esther was launched in June 2008 via www.moddb.com, one of the key mod community sites. At the time of writing, it has been downloaded over 4000 times and feedback has enabled us to assess its reception. This suggests that using mods as research tools has a real value in terms of reaching gamers, outside more formal academic settings. The response has been extremely positive, with negative comments focusing on three things. Firstly, shortfalls in the quality of the environment’s build. Secondly, some players dislike the slow speed and ‘bunny-hop’ to speed the game up, causing multiple triggers to fire over each other; others find ways to escape the maps and are frustrated by clips used to prevent this - a diegetic alternative was created for the second iteration. Finally, and inevitably, others hate the lack of gunplay and find the whole thing boring. What is surprising for us, however, is how few comments of this sort Dear Esther has attracted: generally, the community has loved the mod and it has even

drawn comparisons to commercial games, S.T.A.L.K.E.R. being the most notable example. Attention has been drawn to the quality of the music, voice acting and writing. The last of these is the most surprising as the ambiguity, contradictions and lack of closure in the script is where Dear Esther really deviates from a normal game. However, players seem to understand that this is due to an open, abstract plot rather than structural problems in the delivery; the notion of an unfolding mystery that is never solved seems to appeal to them. Whilst some players find the lack of action frustrating, for most, the atmosphere and drive to find out more about the story is enough of a pull to get them all the way through the experience and for a few to go immediately back and play again for further answers. For these players, the speed of the game is not an issue, and they draw attention to the need to consider Dear Esther as requiring a different attitude to a game system. In terms of affective response, many players report being scared and others describe the experience as eerie, moving and very sad. These last two are emotions that normally fall beyond the affective range of games, especially first-person games.

Dear Esther has proved that FPS players want more than just guns and gore: work that is driven purely by story, with no action or goals, no closure and that denies the player any final understanding of what they have experienced. The value of this kind of storytelling experiment being released to the public domain is clear.

References

1. Mateas, M., Stern, A.: Build It to Understand It: Ludology Meets Narratology in Game Design Space. In: Proceedings of DiGRA 2005 Conference: Changing Views, Worlds in Play, July 16-20 (2005)
2. Valve Half Life 2 Valve, PC (2004)
3. thechineseroom Dear Esther, www.moddb.com/mods/dear-esther (2008) PC GSC: Game World S.T.A.L.K.E.R.: Shadow of Chernobyl THQ / GSC (2007) PC

Walking the Edit – A Research Project of the Master Cinema Network in Switzerland

Ulrich Fischer

ECAL, 5 Av. du Temple
1020 Renens, Switzerland
info@walking-the-edit.net

Abstract. This research project combines the fields of cinema and new media in order to experiment narrative constructions between an existing audiovisual database and a recorded trajectory of a person through an urban space; build through a computed analysis of the path, the result will be an online presented movie.

Keywords: Audio-visual database and metadata, automatic editing engine, interface design, urban space, participative concept, geo-localisation (GPS), cartography.

1 Introduction

The goal is to create documentary movies with narrative qualities out of a database with existing still images (or photographs), videos, texts and sounds. The research project “Walking the Edit” proposes to achieve this with an artistic approach that will lead to an understandable logic for the spectator / visitor: the “path” of the edit will be directly translated from a recorded urban walk.

There will be as much movies as urban walks – a single visitor can create several movies that will be uploaded to a web site where he/she could watch and then share the resulting movies with a larger audience.

1.1 Actual Standing of the Project and Schedule

During the first phase of the project (January – June 2008), we developed a first prototype (software and hardware), a database (based on PostGre SQL), a CMS to edit the media (created specifically for the project) and an editing engine.

The second phase (September 2008 to June 2009) will permit to create enough content to make the concept work, to finalize the software and hardware developments through conceptual and technical experimenting.

Until the end of 2008: improvement of the database, the CMS and the editing engine; testing of the hardware platform; elaboration of the editing concept.

Mid 2009: public presentation of the project (artistic results) and publication of the research conclusions (scientific results).

2 Project Presentation

The project is constructed upon 3 specific moments.

The First Moment is the Creation and the Tagging of the Content. A chosen group of people (cameraman's; filmmakers; editors) elaborate a concept of creating media files (audio, photography, text and video) that reflect different kinds of relationship with a given urban space. Those media files, once created, are then placed within a database and tagged with specific metadata that will permit a narrative combination of those files.

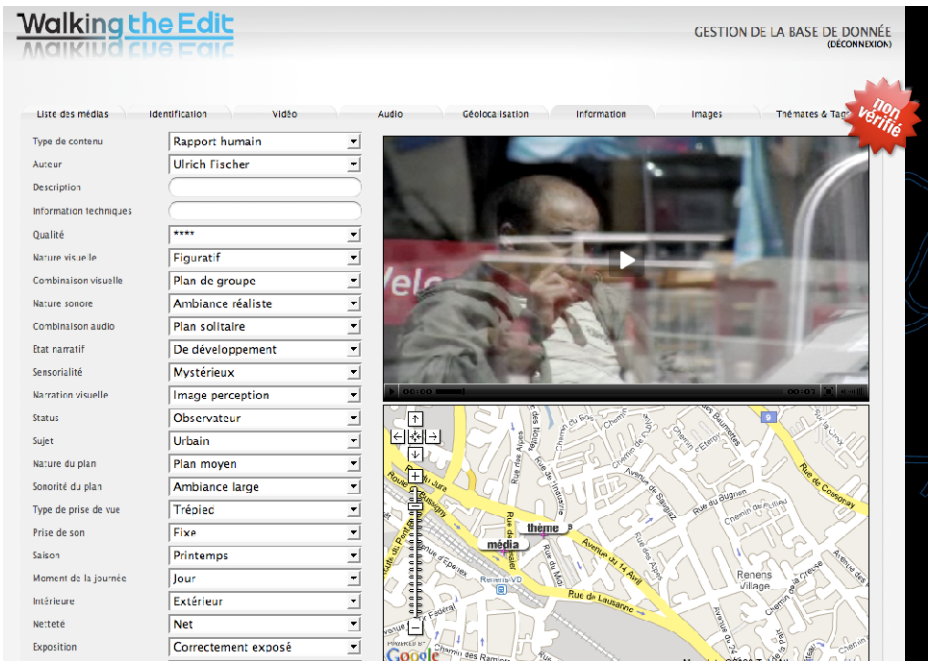


Fig. 1. A view on the CMS

Every media is placed upon a map and receives a unique spatial address; the system will then be able to simulate the combinations of media in order to fine-tune the tagging. The media will have a spatial impact that will be visualized through circles that are calculated through the length of the media; the combination of those circles will in turn create “dataclouds” that will be shown over the real map.

The Second Moment is the Creation of the Real Trajectory and Its Analysis. The project invites a visitor to walk through a part of a city equipped with a special mobile device (with GPS) that gives a view of the city map with a cloud of circles over it (the representation of the media files). For the moment, the type of mobile device is not defined: it could be an Iphone, another smartphone, or a homemade portable device.

The visitor can choose within 3 different clouds that reflect the 3 point of view that where used to create the content: analytic (outside); implicated (inside); poetic (in-side/outside). The choice of the cloud will have a direct impact on the chosen media and therefore on the style and content of the film as it will change the kind of POV.




Fig. 2. A view on the interface of the mobile device

During the walk, a live mix of sound files locally stored will be streamed to the visitor – this sound composition will give a suggestive idea of the film that is currently being walked. This means that the editing engine (and a sound database) is also build into the mobile device. The visitor can also see it's current position and the trace of the trajectory on the map (as shown in the example above); the only way to influence the result is to set the In and Out point of the movie (begin and end of recording), to set the type of cloud, and finally to decide how and where to walk.

Once the walk is done and the device given back, the analysis engine placed on a media server will create a play list out of the raw GPS data and some predefined rules that are implemented in the database. The precision of the analysis made at this point will be higher than the real-time analysis, as the whole process can begin from the end and therefore produce better combinations as during the walk.

The whole editing process must be all automatic; to achieve this, it must be highly customizable in order to fine-tune the way the combinations can be done. All the software developments are done with open source software under respect of current standards.

The Third Moment will Present the Result: The Movie(s). Once the edit done, the system automatically transfers on a website the resulting video file to the visitor (who now becomes a spectator).

Fig. 3. A view on the web interface with the film and the path

It will be possible then for the walker to view his / her own movie on a private page, with possible actions and interactions: to scroll the movie on the map by changing the place of the position of the visitor (on a sort of “space-line”), to view statistics of the trajectory (visualization of the metadata combined with the analysis of the path). On this page will also be the other films created by other visitors and some general information on the project.

Once the person is ok with it’s walked and manipulated edit, it will be possible to publish the video on a public page where all the other movies are accessible to the public.

Reference

1. Fischer, U.: Walking the Edit. Project homepage (Last accessed 18.09.2008), <http://walking-the-edit.net/>

3D Immersion in Virtual Agents Education

Cyril Brom, Jakub Gemrot, Ondřej Burkert, Rudolf Kadlec, and Michal Bída

Charles University, Faculty of Mathematics and Physics
Malostranské nám. 2/25, Prague, Czech Republic
brom@ksvi.mff.cuni.cz,
{jakub.gemrot,ondrej.burkert,rudolf.kadlec,
michal.bida}@gmail.com

Abstract. Many projects featuring intelligent virtual agents have emerged in last years, but not many reports on their advances in education. This paper presents the curricula of a university course on Modelling Behaviour of Human and Animal-like Agents, including a seminar in which students develop their own virtual agents using a toolkit we have developed. This course has been also scaled for a workshop with computer science high-school students. An evaluation of the course is presented and main lessons learned overviewed. The paper also explicitly formulates the teaching methodology underpinning the course and outlines several general questions hoping to start a broader discussion on educational issues related to the field of intelligent virtual agents.

Keywords: Intelligent virtual agents, storytelling, education, course curriculum.

1 Introduction

The field of intelligent virtual agents (IVAs) and interactive storytelling (IS) is becoming matured. As the community is growing, an issue how to educate its new members, in most cases undergraduate students, starts to demand attention. Educational issues seem to be well addressed by neighbouring disciplines, namely computer graphics, robotics, and software agents; consider e.g. the number of summer schools and introductory books available for these domains, the RoboCup platform, or the special educational track at the Eurographics conference. While the classes directly focused on IVAs have already started to emerge, and some of the class materials are available on-line, e.g. [2, 6], the reports on advances in education directly related to IVAs and IS are scarce.

In this paper, we report on two things. First, the paper describes the curricula of a one-term theoretical course on *Modelling Behaviour of Human and Animal-like Agents* [3], highlighting main lessons learned from teaching perspective. The course was created in 2005 primarily for undergraduate computer science students, but it was also adopted for a high-school workshop and for a summer school. While primarily focussing on modelling behaviour of IVAs in the context of interactive applications such as serious games and virtual storytelling, the course also provides students with limited knowledge of cognitive and behavioural sciences, including emotional, perception, and memory theories, natural neural networks, and basics of ethology, presenting a starting point for students considering studying these subjects in the future. As far as we know, there is no single book covering the whole curricula.

Second, the paper describes integration of this theoretical course with a practical seminar, developed in 2008, during which students learn to program their own virtual characters. Although there is an excellent entry-level tool for building simple agents and running social simulations that is used for seminars concerning classical software agents, NetLogo [19], we do not find it suitable as a main tool for students attending an IVA/IS course. It does not allow for creating agents with complex behaviour and possesses no complex 3D environment. This prevents students from exploring advanced applications such as virtual storytelling or 3D games, which, in effect, reduces their motivation. Computer game development courses tend to capitalise on 3D game engines like Counter Strike, Quake or Unreal Tournament (UT). However, coding the behaviour of a character directly in the language of a game is cumbersome for beginners. They must focus on low-level issues related to communication, navigation etc.; consequently, after dozens of hours of work, their agents "still don't do anything cool", which can be quite discouraging. In our seminar, students develop their characters for Unreal Tournament 2004 (UT04) [7], but instead of using Unreal Script, the native scripting language of the UT04, they use a special purpose toolkit Pogamut 2 [12] connected to UT04 via a Gamebot-like [1] interface. Pogamut 2 was developed to facilitate start with IVAs, both for educational and research purposes, and it is freely available to download¹. Importantly, Pogamut 2 allows for rapid development of simple "NetLogo level of complexity" agents as well as challenging, complex agents with deliberative capabilities, thereby supporting a two-stage process of education: "doing simple things quickly" as well as "exploring the more advanced stuff".

The primary goal of this paper is to discuss the curriculum of the course in order to facilitate either development of a similar course or adoption of our toolkit for a course that is already running. In this respect, the paper is most similar to the work [8], which reported on augmentation of a course on *Autonomous Agents and Multi-agent Systems* with a practical seminar using Counter Strike 3D game. We depart from this work in that our course is directly focused on IVAs behaviour, and we use a toolkit featuring not only a multi-agent library but also an integrated development environment tailored to support the development of IVAs behaviour. The number of supplementary features of Pogamut 2 and its support of education distinguishes it from similar freely available tools, including solutions capitalizing on Gamebots [1], as discussed in depth in [12]. On a more general level, the aim of this paper is to start a broader discussion on educational issues related to the subject of IVAs and IS.

The paper proceeds as follows: Sec. 2 details the course's curricula, including main lessons learned concerning individual course's parts. Sec. 3 introduces Pogamut 2 and explains how it is used in the seminar. Also a small evaluation study is presented. Sec. 4 discusses a teaching methodology underpinning this course, several general lessons learned, possibility of transferring the course to another teaching context, and also opens the general discussion on education of the subject of IVAs and IS.

2 The Course

The course was created at Charles University in Prague, the computer science study program, in 2005. Its theoretical part comprises 13 lesson units (1 unit = 90 minutes).

¹ The toolkit is available for download at <http://artemis.ms.mff.cuni.cz/pogamut>

In 2008, the practical seminar has been added (6 units). At the end, each student must create his/her own IVA, which presents about half of the student's evaluation.

The course is tailored to computer science students at least in their fourth term of bachelor studies, after they attended several courses on programming (10.5²), mathematics (17), general IT skills (8.5), and algorithms (5.5). Most of them have only limited knowledge of other topics. Many are recruited from technical high-schools. Every year, the course is attended by about 30 – 50 students (about 15% of the total number of students of one grade).

Objectives of the course are (1) to introduce the field of interactive applications featuring IVAs, (2) to teach students to develop behaviour of IVAs, (3) to boost their interest in related disciplines, namely in artificial intelligence, autonomous agents, and behavioural sciences. There are many other courses where the students can learn about other aspects of IVAs, most importantly courses on computer graphics, computational linguistic, artificial intelligence (AI), autonomous agents (AA), and computer games development (see [3] for details). While there is no given order in which to attend these courses, the questionnaire administered this year showed that our lecture is the entry-level point for the subjects of AI and AA for 60% students (either the first or the only lecture related to AI). Computer graphics, on the other hand, is studied mostly earlier than or in parallel to our course. From the point of view of IS and IVAs, some courses that would give students a broader context are missing; most notably social sciences and human-computer interaction.

Curriculum. Conceptually, the course comprises an introductory lecture, three theoretical blocks, and the practical seminar. The curriculum of the theoretical part is overviewed in Tab. 1. In general, the course starts with relatively concrete models for controlling behaviour and proceeds to more abstract ones. The course has more possible orderings of lectures, specifically within the block on advanced topics (Fig. 1). The content of the course is detailed in the supplementary material to this paper [3]. The rest of this section highlights main lessons learned. The motivation for the gross structure of the course is discussed in Sec. 4.




Fig. 1. Possible ordering of lecture parts

Introduction. The knowledge of many students on IVAs is minimal. They typically think that IVAs belong only to the domain of computer games. Thence, it is crucial to demonstrate the broad scope of the field. Students most welcome live state-of-art demonstrations. In this way, interactive drama, serious games, cultural heritage and

² Normalised number of courses on a given topic; "1 course" equals 13 lesson units (i.e. 13 x 90 minutes). The course presented in this paper, with the practical seminar, amounts to 1.5.

film applications, cognitive science research and therapeutic applications, technical applications, and computational ethology should be introduced. The last two examples are important for drawing distinction between *plausibility* in the sense of a natural science, e.g. physics or ethology, and *believability*. The NetLogo examples help with clarifying the notion of *emergence*. The notion of *autonomy* has to be stressed. The survey of main issues, such as navigation, action selection, emotional modelling, story generation, helps to outline the content of the rest of the course.

Table 1. The curriculum of the course

Block	Topic	Units
	(1) Introduction	1
Concrete models of action-selection.	(2) Reactive planning: hierarchical if then rules with priorities, hierarchical state machines, fuzzy rules, probabilistic state machines.	2
	(3) Neural networks, reinforcement learning, evolutionary algorithms.	2
	(4) Behavioural science models: "Psycho-hydraulic" model of K. Lorenz, Tyrrell's free-flow hierarchy, classical and operant conditioning.	1
	(5) Path-finding, steering, abstract terrain representations.	1
Conceptual notions, architectures, representation.	(6) Architectures: symbolism vs. connectionism, layered architectures, notion of deliberation. BDI, planning.	1
	(7) Multi-agent systems introduction: types of agents, communication.	1
	(8) Representation: logic, deictic repres., Gibson's affordances, smart objects.	0.5
Extras, broader context	(9) Artificial emotions.	1
	(10) Storytelling: emergent narrative, plots representation, level-of detail AI.	1
	(11) Perception and memory.	0.5
	(12) Unified theories of cognition.	1

Concrete models of action selection. This block helps students to see that some action selection issues can be promptly solved and how. Students appreciate concrete mechanisms and concrete solutions. Our experience is that they tend to regard these special-purpose mechanisms as abstract and general. While it is useful to give concrete counter-examples in which these methods fail, it is ineffective to give students any abstract, background knowledge, at this moment (more on this in Sec. 4). In fact, these concrete mechanisms will serve as a base upon which the students will represent in their minds more advanced and/or general knowledge.

An excellent material for introducing neural networks and evolutionary mechanism is the control architecture of animals from the computer game *Creatures* [10]. First, it helps to explain both topics. Second, the neural networks used in *Creatures* are more biologically plausible than typical artificial neural networks, which helps with introducing the distinction between biological and artificial neurons. It is also useful to introduce the action selection problem from the perspective of behavioural sciences. This demonstrates rich cross-fertilisation between disciplines, and reminds students of the believability—plausibility distinction. The excellent material is [18].

The lecture on path-finding helps with two additional things. First, the notion of representations can be stressed (with the help of terrain representations). Second, when drawing the distinction between low-level steering and high-level path-planning, it will become apparent that it can be helpful to think about IVAs at different levels of analysis - a bridge to the next part of the course.

Concept, architectures. Technically, so far, the students have been taught special-purpose solutions. Now, their knowledge should be unified within more abstract, conceptual frameworks. Notions of goals and intentions, basically Belief-Desire-Intention (BDI), should be introduced, reactive approach confronted with planning, symbolism with connectionism. A welcomed example demonstrating that BDI can be really implemented in various ways is the hybrid control architecture of Black & White creatures [9]. To frame the notion of IVAs in a broader context and to discuss the topic of communication, it is vital to introduce multi-agent systems.

Extras. Every year, many students are surprised that when developing an IVA, despite it is intended to be believable but not psychologically plausible, it might be a good idea to look at what psychologists say about real humans. For example, the students are amazed that mechanisms of attention can really help with limited computational resources. In sum, these lectures help the students to understand that a "mind" of an IVA comprises not only an action selection mechanism, but also various additional "circuits". Our experience is that it is better to give these lectures after the students are already familiar with IVAs architectures and action selection.

The high-point of the course is storytelling for many reasons. Most importantly, students start to understand that IVAs cannot exist without virtual worlds and vice versa, in other words that IVAs and virtual worlds constitute an intertwined couple, a sort of *marionette theatre* in a service of a user. Through this metaphor and the notion of *story manager* students can understand how to *systematically* abandon the concept of strong autonomy. These notions also help with introducing the technique of level-of-detail AI [4]. Additionally, the notion of *emergent narrative* reminds students of the idea of emergence and the notion of *story construction* reminds them of planning.

3 The Seminar and the Toolkit Pogamut 2

The aim of the practical seminar is to give students an opportunity to exploit theoretical skills acquired during the course, i.e. to implement own IVA on a chosen topic, such as state machines or artificial emotions. As the course covers many topics, we looked for a toolkit having a potential to provide a base for as many topics as possible while not being too complex and becoming a burden to students. Such a solution wasn't available in previous years [12]. Therefore we have created the toolkit Pogamut 2.

Features of Pogamut 2. Pogamut 2 is designated for educational and research projects concerning modelling of behaviour of IVAs. It brings together a 3D simulator, an agent library, and a development environment (IDE). The toolkit uses the environment of Unreal Tournament 2004 (UT04), an action game as the 3D simulator. The UT04 offers many maps of complex 3D worlds as well as human-like models of agents suitable for games and storytelling. The agent library, written in Java, allows students to quickly create new agents inside the UT04. It also enables utilisation of third-party libraries and planners (such as neural networks or POSH [5], which is already integrated). Rapid exploration is further facilitated by the IDE, which is based on the popular NetBeans platform [17]. It helps newcomers to get quickly comfortable with the process of agent development, namely:

- It offers agent project management tool along with several example projects.
- It has two types of view – Simple and Advanced. The simple view is used in the first lesson, allowing for quickly setting up and running example projects to motivate students. The latter one is used during the agent development.
- It gives contextual help during coding of the agent eliminating the necessity for consulting the manual frequently.

Additionally, the toolkit possesses many advanced features, such as a debugger, a variable manager, a simplified map of the environment facilitating orientation of the user, a support for GRID computing, etc. [12].

Table 2. Curriculum of the seminar. The two stages are separated by double line.

(1) Introduction, example projects	The toolkit and the sample IVAs were introduced. Students interactively explored the examples, meaning they made simple changes of the IVAs behaviour.
(2) Implementation of a simple IVA	Students were taught the agent library basics. Finally, they were able to implement their own agent, which was able to follow the player inside the UT04. Students were delighted that such behaviour can be achieved with merely 4 lines of code.
(3) Navigation	The topics: How to navigate in UT04 using way-points. How to detect obstacles using ray-casting. Finally, students were able to create an agent running around the environment, collecting weapons along the way, and performing dodges.
(4) Reactive planner POSH	Students learned how to code POSH-behaviour. Finally, they were able to reproduce several sample agents using POSH-rules.
(5) Items, weapons and weapon handling.	Students learned how to handle different weapons. Finally they were able to produce agent with basic combat capabilities. (Note, that the problem of weapon selection and usage is an issue of its own studied in gaming AI)
(6) Practise	This lesson was reserved for student's projects.

Seminar topics. The seminar comprises six lessons at computers (Tab. 2). It capitalises on a two-stages teaching process: rapid start with simple things and slower exploration of more advanced stuff. Note, that the topics of the basic stage can be actually covered by NetLogo, contrary to the advanced stage.

Since the toolkit was introduced for the first time, we decided to keep its utilisation as simple as possible. Actually, we exploited the fact that UT04 is the underlying platform and focused mainly on gaming AI issues. Since we believe that these present only a part of what should be explored by students in such a seminar, we are presently adding a generic emotional module and a story manager to the Pogamut 2 library. These are intended to become the topics of the 6th lecture in the next year.

Concerning students' assessment, each student had to develop his/her own IVA. The default goal was to create an IVA for a death-match tournament. Nevertheless, many students came up with their own tasks. Ideas ranged from cooperating agents that would support themselves in the death-match over explorer agents building their own environmental representation to evolution of a neural network controller that would use ray-casting alone to navigate through the environment. As students were

enthusiastic about their own ideas, it worked best to let them explore those ideas on their own even if their goal has apparently been too complex.

Evaluation. In order to evaluate the seminar and our approach, a pre- and post-questionnaires were given to 25 of our students (M=23, F=2), corresponding to 55% of total enrolled students. For brevity, only the main findings are highlighted.

The students were required to provide free report on their previous background, their motivation, and whether they think the course had helped them to understand the problematics. Table 3 summarises these data. It is apparent that the course was attended by students with various levels of previous knowledge both in AI and programming. Nevertheless, we found no correlation between this stratification and other reported data. Most students attended the course out of curiosity. This is an interesting finding for it suggests that the subject of IVAs/IS is still new and there is only limited previous common knowledge about it among students. Also it seems that not many students believe they can earn money by studying this topic (but can they?). Finally, the table shows that the students regarded the seminar beneficial.

In order to evaluate the platform, the students had to mark its features from 1 (worst) to 5 (best). The scores are high, except for the documentation, which we have started to improve.

Table 3. Summary of the main findings from the questionnaires

Question:	Category:	Nr.
Previous background in AI / programming?	Few (AI: 0 – 2 courses [normalised] / Prog.: 2 years studying)	15 / 8
	Some (AI: 3 – 5 courses / Prog.: 3 – 4 years studying, an internship)	5 / 4
	A lot (AI: 6 – 12 courses / Prog.: more years studying, at least several months of IT employment)	5 / 13
Why did you choose this seminar? (more answers are allowed)	Out of curiosity	21
	Fond of UT04	4
	I want to study Gaming AI / Entertainment applications	3
	Sounds like a good practical seminar	2
	To see AI in practise	1
Did the seminar help you to deepen your knowledge about the action selection mechanisms of IVAs?	Yes, it helped.	9
	It did not deepen the theoretical insight but it gave examples of practical problems the ASM must address.	11
	No, it didn't.	1
	Can't say.	4

Table 4. Evaluation of the features of the platform and comprehensiveness of the course

Easy to learn	3,82
Documentation	2,45
GUI and coding user-friendliness	3,91
NetBeans environment	3,64
Comprehensiveness of the theoretical course	3,82
Comprehensiveness of the practical seminar	4

Finally, we asked the question "Do you find Pogamut 2 toolkit suitable for practical needs of the course or should we look for something else?" Students assessed

Pogamut 2 as quite appropriate for the demonstration of practical implementation of IVAs behaviour. They really appreciated the fact that simple things can be expressed with minimum code, which encouraged them to experiment with the toolkit. Some students appreciated the UT2004 game and even suggested to connect Pogamut 2 with different environments as well. Others, especially the girls, demanded a non-violent environment. Examples of concrete reports follow:

"Pogamut 2 is the right tool for the course."

"Simple things in simple way, complex possible."

"It would be cool to connect Pogamut 2 with Operation Flashpoint to provide realistic maps and to try to implement some complex behaviour according to military books."

"Some realistic environment allowing better storytelling application would be nice."

We believe the emotional module and the story manager will help with the last point. We also look for another environment to connect the toolkit with, the Second Life [14] is one of the possibilities we are considering presently.

Case study. Adoption of the course for high-school students. An important feature of every course is whether it can be transferred to another learning context. In general, the issue of scaling of our course is discussed in Sec. 4. To back up this debate, it should be said that we have already transferred the course. Its theoretical part was adopted for a summer school (four 90 minutes lessons). The whole course was also scaled for high-school students of computer science (16, 17 yrs old, 2nd grade). This class lasted 90 minutes of theory and 90 minutes of practice with our tool. Despite the limited time, the students were given a compact introduction into the subject. The issues of path-finding, steering and emotions were discussed, and the links towards ethology were highlighted with the help of NetLogo. Navigation issues were solved with the usage of reactive planners and simple neural networks controllers thus giving concrete solutions for concrete issues. Finally, a brief introduction into Pogamut 2 was made, during which the students had an opportunity to implement a simple IVAs behaviour by means of a few if-then rules. The questionnaires were administered; the results are similar to the results of university students (Tab. 3, 4).

4 General Discussion

Many courses on computer games development and multi-agents systems capitalises on the same teaching methodology as our course. The objective of this section is to make this methodology explicit. This will help us to discuss structure of such courses in general and also highlight several findings concerning our course that have not been discussed yet.

Let us perceive the course as a learning process. Now, we can say that we engaged our students in such a process, in which they could systematically explore the field of study, incrementally constructing their personal representations of knowledge, that is, a) objectives of the fields of IVAs/IS: motivations, goals, applications, b) technical issues stemming from these objectives, c) models/algorithms/mechanisms of IVAs and virtual environments that cope with these issues, i.e. solutions (Fig. 2). Students were not only to mentally represent this knowledge, but also to develop new mental

processes to operate with this knowledge. Their general ability to think about algorithms and models from different perspectives was intended to increase. Importantly, this ability is quite general, applicable *outside* the field of IVAs.

What were the stages of this learning process? The retrospective analysis suggests that they were three, roughly corresponding to the blocks of the course, helping the students to build their knowledge *incrementally* (Fig. 3). First, the students were confronted with different technical issues of the IVAs and several "middle-level abstraction" mechanisms that can solve these issues, such as prioritised if-then-rules, finite state machines, reinforcement learning, steering rules etc. Second, they were to understand that these models are embedded within general, let us say conceptual or "black-box", frameworks, such as rule-based systems, first-order logic, or BDI. They were required to "zoom through" different levels of description and be able to mentally substitute one model for another. For example, they had to "see" that neural networks controlling animals from the game *Creatures* [10] are conceptually similar to fuzzy if-then rules. As computer science students, they also had to acquire the ability to bring their mental models to "life" in the digital media, i.e. to be able to "zoom in" at the code level of description. At the same time, they had to start to understand that the whole point of this enterprise is not the issues, but the objectives. The third stage even strengthened this notion by means of providing the students with additional objectives, issues and solutions, and links with other disciplines to help them to connect their new knowledge with what they have already known and with what they might learn in the future. To help the students to create this "semantic network" (Fig. 3(3)) was the primary objective of the course.




Fig. 2. The conceptualisation of the causalities behind the course

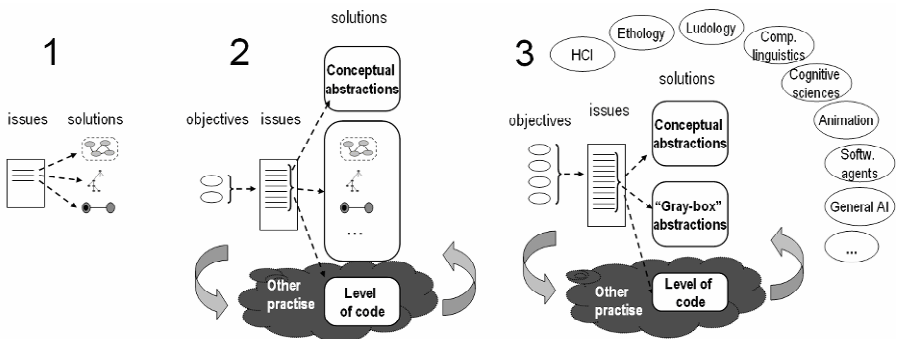


Fig. 3. Schematic depiction of knowledge students are intended to represent within their minds during the three stages of the learning process

This would not have been possible, however, without providing *feedback*, through which the students could test their ideas. This feedback was given via two mechanisms: through the practical seminar and by demonstrating various applications during the course, which the students should have further explored during self-studying (e.g. Façade, NetLogo).

This teaching framework is actually situated within the *constructivist* learning approach [15]. The structure of knowledge the students should build can be conceived as a kind of mental models of Johnson-Laird [11]. It is natural to suggest a metaphor of a *non-linear game* for courses capitalising on this framework. The non-linear aspect stresses that students should be allowed to approach the field from *different perspectives*, visiting the topics in order and at speed given by their individual preferences and background knowledge. The *gaming* aspect stresses that the process of exploration should be *interactive*.

Now, what does this tell us? We will start with concrete points and proceed to more general ones.

1. Bottom-up approach. The described teaching process is a bottom-up strategy: to start with a concrete and to end up with an abstract. One year we swapped the first two stages depicted in Fig. 3 and the interviews showed that this was not a good idea: the students forgot or did not grasp the abstract without the concrete at all.
2. Transfer of our course to another learning context. Basically, the course can be scaled along two axes: (A) background knowledge of the students and other courses available, (B) available time. Along both lines, the following key points should be preserved: the availability of feedback and incremental knowledge development.
 - (A) We have three suggestions. (i) To change the feedback according to previous students' knowledge. Practise of our students is focused on programming, but for non-IT students, the advanced programming block (see Sec. 3) can be replaced e.g. by developing Machinimas [13] or building virtual worlds in Unreal Editor. (ii) Links to other science disciplines should anticipate what knowledge students are most likely to need in the future. (iii) While some concrete topics can be suppressed in Stage 1, e.g. neural networks, some extras can be added in Stage 3, e.g. agents' social interaction.
 - (B) Again, we have three suggestions. (i) In Stage 2, building knowledge of conceptual models can be suppressed. Our experience is that without enough time available, the students are not able to grasp the abstract knowledge anyway. (ii) In Stage 1 and 3, many concrete issues and solutions can be omitted (e.g. probabilistic state machines, Tyrrell's free-flow hierarchy, perception and memory). Nevertheless, it is important to preserve the line "concrete issues → concrete solutions" during the Stage 1 at least to some extent to boost up students curiosity and engagement. (iii) Discussions about connections with other sciences can be shortened, but again, at least some links should be provided.

Along these lines, the course was actually scaled for a summer school and a high-school course, as detailed in Sec. 3. Many points of this "transfer schema" are quite general and can work for similar courses as well.
3. Availability of materials. Today, predominantly, research prototypes, applications, or at least videos, are *not* published on the internet by their authors. As these materials are essential for "providing feedback", this trend should be changed. This "call for teaching resources" is even amplified when considering

correlations between availability of materials and community growth (think of open-source communities or communities of gamers emerging around partially-opened games such as Unreal Tournament or Counter Strike).

4. Another open issue is a schoolbook that would cover the topics of the fields of IVAs/IS. Though there are collections of papers, e.g. [16], these cannot compensate for a coherent learning text. Neither can this be done by general AI, gaming AI, software agents or computer graphics textbooks. We argue that interactive applications featuring IVAs is a subject of its own and requires its own schoolbook covering basics ranging from animation through action selection and virtual storytelling to social and behavioural sciences, giving the reader a consistent and a broad view on the field.
5. Longitudinal evaluation. Currently, there are only very limited data available on how many students of interactive technologies actually use their knowledge in practise. Similarly, it is not yet known how many students really keep in their minds the knowledge they gained during curricular learning: longitudinal studies are missing. E.g. does the proposed 3-staged teaching process increase retention of knowledge over long periods? How many students earn money doing interactive applications comparing to let's say databases? We have to start to ask such questions.

5 Conclusion and Future Work

This paper started with drawing attention to education of new members of IVAs and IS community, an issue that has not been much studied yet. From this perspective, the contribution of this text is twofold: first, the curricula of a university course on *Modelling Behaviour of Human and Animal-like Agents* has been overviewed and the main lessons learned (from the educational perspective) detailed. Second, the teaching methodology behind this course and similar courses has been verbalised to facilitate thinking about such courses and about their transfer between different learning contexts. A small evaluation study has suggested that the course is effective and a case-study demonstrated that the transfer of the course is possible. Nevertheless, a rigor longitudinal evaluation is missing.

The paper also introduced the toolkit Pogamut 2, a training vehicle in which students can practice development of IVAs using the environment of Unreal Tournament 2004. Actually, the toolkit can be used for research purposes as well.

Presently, the Pogamut 2 is being augmented with a generic emotional module and a story manager, two components that we believe will contribute most to the learning experience of students. We are also considering connecting Pogamut 2 with another environment, e.g. Second Life [14].

In general, the authors hope that this paper will help to start a broader discussion on educational issues related to the subject of interactive storytelling and intelligent virtual agents, including such topics as learning materials, a "required minimum" for a student of this field, or facilitation of team work among students with different backgrounds (e.g. programmers and artists).

The lecture materials, the extended version of this paper detailing the curricula and the literature used, and the Pogamut 2 are freely available for download [3, 12].

Acknowledgments. This work was partially supported by the Program "Information Society" under project IET100300517, by the Ministry of Education of the Czech

Republic (Res. Project MSM0021620838), and by two student grants GA UK No. 351/2006/A-INF/MFF, and 1053/2007/A-INF/MFF.

References

1. Adobbati, R., Marshall, A.N., Scholer, A., Tejada, S.: Gamebots: A 3d virtual world test-bed for multi-agent research. In: Proc. 2nd Int. Workshop on Infrastructure for Agents, MAS, and Scalable MAS (2001) (26.6.2008), <http://gamebots.planetunreal.gamespy.com/>
2. Aylett, R.: Slides for the course on Artificial Intelligence, Heriot-Watt University, Edinburgh (26.6.2008), <http://www.macs.hw.ac.uk/modules/F23AD3/AIcourse.html>
3. Brom, C.: Supporting material to this paper, extended version of this paper, slides for the course on Modelling Behaviour of Human and Animal-like Agents. Charles University in Prague (26.6.2008), <http://artemis.ms.mff.cuni.cz/main/tiki-index.php?page=Teaching>
4. Brom, C., Šerý, O., Poch, T.: Simulation Level-of-detail for Virtual Humans. In: Pelachaud, C., Martin, J.-C., André, E., Chollet, G., Karpouzis, K., Pelé, D. (eds.) IVA 2007. LNCS (LNAI), vol. 4722, pp. 1–14. Springer, Heidelberg (2007)
5. Bryson, J.J.: Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents. PhD thesis, MIT (2001)
6. Dignum, F., Westra, J.: Slides for the course on Games and Agents, University Utrecht (26.6.2008), <http://www.cs.uu.nl/docs/vakken/gag/>
7. Epic Megagames: Unreal Tournament (2004) (26.6.2008), <http://www.unrealtournament.com/>
8. de Melo, C., Prada, R., Raimundo, G., Pardal, J.P., Pinto, H.S., Paiva, A.: Mainstream Games in the Multi-agent Classroom. In: Proc. Int. Conf. on Intelligent Agent Technology (2006) (26.6.2008), <https://fenix.ist.utl.pt/disciplinas/aasm>
9. Evans, R.: Varieties of Learning. In: Rabin, S. (ed.) AI Game Programming Wisdom I, pp. 567–579. Charles River Media, Inc., Hingham (2002)
10. Grand, S., Cliff, D., Malhotra, A.: Creatures: Artificial life autonomous software-agents for home entertainment. In: Proc. 1st Int. Conf. Auton. Agents, pp. 22–29. ACM press, New York (1997)
11. Johnson-Laird, P.N.: Mental Models. In: Posner (ed.) Foundations of Cognitive Science, ch. 12. MIT Press, Cambridge (1993)
12. Kadlec, R., Gemrot, J., Burkert, O., Bída, M., Havlíček, J., Brom, C.: Pogamut 2 – a platform for fast development of virtual agents' behaviour. In: Proc. CGAMES 2007, La Rochelle, France (2007) (26.6.2008), <http://artemis.ms.mff.cuni.cz/pogamut>
13. Kirschner, F.: Movie Sand BOX (26.6.2008), <http://www.moviesandbox.net/>
14. Linden Research, Inc.: Second Life (26.6.2008), <http://secondlife.com/>
15. Pappert, S.: Mindstorms: Children, Computers, and Powerful Ideas. Basic Books, New York (1980)
16. Prendinger, H., Ishizuka, M. (eds.): Life-Like Characters: Tools, Affective Functions, and Applications. Springer, Heidelberg (2004)
17. Sun Microsystems, Inc.: Netbeans (26.6.2008), <http://www.netbeans.org>
18. Tyrrell, T.: Computational Mechanisms for Action Selection. PhD thesis, Centre for Cognitive Science, University of Edinburgh (1993)
19. Wilensky, U.: NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University (1999) (26.6.2008), <http://ccl.northwestern.edu/netlogo/>

Exploring Non-verbal Behavior Models for Believable Characters

Magy Seif El-Nasr and Huaxin Wei

Simon Fraser University, School of Interactive Arts and Technology,
Surrey, BC, Canada
{magy, huaxinw}@sfu.ca

Abstract. Believable characters constitute an important component of interactive stories. It is, therefore, not surprising to see much research focusing on developing algorithms that enhance character believability within interactive experiences, such as games, interactive narrative, and training environments. These efforts target a variety of problems, including portraying and synchronizing gestures with speech, developing animation tools that allow artists to manipulate and blend motions, or embed emotions within virtual character models. There has been very little research, however, devoted to the study of non-verbal behaviors, specifically mannerisms, patterns of movement including postures, gaze, and timing, and how they vary as a function of character attributes. This paper presents a work in progress of a study conducted to (1) identify key character characteristics recognized by animators using an acting model, and (2) formalize non-verbal behaviors patterns that animators use to express these character characteristics.

Keywords: Believable characters, animation, acting, virtual characters, embodied agents, articulate 3D characters.

1 Introduction

Believable characters play an important role in many interactive entertainment productions, including computer and video games, training simulations, and educational games [1, 2]. Current industry methods rely on heavy scripting, where voice acting, dialogue scripts, hand-coded animation routines, and hard-coded behaviors are used to portray the desired character; examples of games that employ very detailed motion-captured characters, include *Assassins' Creed* and *Prince of Persia* (developed by Ubisoft) and *Façade* (developed by Mateas and Stern). In these games, artists work very diligently to detail characters' mannerisms and body motion to exhibit the right culture and character characteristics [3]. Such attention to detail of the non-verbal character behaviors is a crucial element for character believability [4]. However, this kind of scripting is labour intensive and rigid, as it does not adapt to all variations induced by interaction.

An alternative is to use artificial intelligent algorithms and graphics techniques to adapt character behaviors to variations in context induced by interaction. This

alternative, however, is not as simple as it sounds, as it has been under research for many years and is still an open problem. Researchers have been working on several fronts to create believable expressive characters that can dynamically adapt within interactive narratives. Graphics researchers, for example, focus on embedding emotions and personality as parameters that can be used to modify virtual character animations [5-7]. Conversational agents researchers focus on building articulate virtual characters that can automatically synchronize gesture and speech [8]. Artificial intelligence researchers focus on integrating models of emotion and personality to build characters that have the ability to improvise [9-11].

As researchers tackle different aspects of this open problem, gaps between these different directions start to appear. One important gap is the gap between character models (artificial intelligence) and how these characters are portrayed through animation (graphics). While there are models that formalize emotional expression through facial muscles [12], there is very little work that explores methods of formalizing non-verbal body motions as a function of character characteristics. We note two previous attempts that looked at body movements as a function of emotions [13] and [14]. These studies, however, focused on emotions rather than character characteristics, such as personality, age, or culture.

In this paper, we study two concepts: *non-verbal behavior patterns* and their relation to *character attributes*. We define non-verbal behavior patterns as: a list of two or more movements linked with specific timing and pacing constraints. For example, the motion of quickly glancing at a character then at the ground is considered a non-verbal behavior pattern. We use the terms character attributes, character characteristics, and character model to mean a list of parameters that define a character, including age, physique, personality, behavior tendencies, quirks, habits, mind-set, and belief system. These concepts are not formalized; our goal is to formalize and define these concepts as part of our ongoing research. In this paper then, we attempt to describe a work in progress exploring two main research goals: (a) develop a set of character attributes that can be used to describe the essence of a character from a narrative perspective, and (b) identify non-verbal behavior patterns that are linked to the character attributes identified in (a).

2 Previous Work

2.1 Believable Characters

The topic of believable characters has been under research for many years. The Oz project presented an early work that developed believable agents for interactive drama [15]. They developed an authoring language for encoding character attributes, such as emotions, personality, and attitudes [16]. They also proposed an agent architecture composed of a reactive planning system which was used to select behaviors, from an authored set of behaviors, dynamically based on context. Mateas and Stern later extended their system by developing ABL (A Behavior Language), which was used to encode behaviors for the interactive drama *Façade*. ABL extended previous work by integrating a mechanism for handling joint behaviors [17]. While *Façade* and the Oz project showed expressive characters that employ several non-verbal behavior

patterns, such behaviors were hand coded by the authors within the authored behavior routines. Thus, there were no formal models used. In addition, several researchers explored the integration of emotions and personality as character attributes within believable characters [10, 11, 18, 19]. While these architectures presented an adaptive routine for selecting behaviors that depend on characters' emotions and attitudes, non-verbal behaviors were manually encoded within the behavior specification. This limits the design as authors still need to hand-code all non-verbal behavior patterns and vary them based on variations in the character models.

There are several graphics researchers who have attempted to address this problem from the graphics end. Specifically, they focus on developing real-time algorithms that modify animation routines, such as walk, run, jump, by adding mannerisms, emotions, and personality [5-7]. For example, Perlin created a framework for *procedural emotion shaders* [20, 21]. The goal of his work is to allow designers to dynamically encode mannerisms for their character animations, and thus they can convey mood, emotions, and very simple personalities through the base movements and actions the animators create. His work has been integrated into Poser and the Half Life engine. Thus, artists can create several variations to their animation by simply selecting an option to modify the animation in a certain way. An example is adding 'sexy' modification for a 'walk' animation developed by the animator. Allbeck et al. developed a similar system for encoding mannerisms in animation [7]. They proposed PAR (Parameterized Action Representation), an action encoding method based on the Laban movement notation¹ [22]. While the examples discussed above have demonstrated great efforts in varying character mannerisms and expressive abilities, they do not address the concept of non-verbal behavior patterns, i.e. including sequence or parallel behaviors with timing and spatial constraints, or relate such patterns to characters attributes, other than emotions and moods.

2.2 Understanding Non-verbal Body Motion Patterns

There are few research projects that attempted to understand non-verbal behavior patterns and their link to character attributes. Wallbott and Scherer [13] presented a seminal work in this area. They studied a sample of 224 videos, in which actors portrayed a variety of emotions in a scenario. Through this study, they found that some body movements and postures can be specifically mapped to certain emotions. For example, 'arms crossed in front of chest' is typical for pride.

Marsella et al.'s work presented yet another example of a study focused on understanding non-verbal behaviors. In their work, they verified Delsarte's model, specifically hand movements [23]. Delsarte was a 19th century musician who developed an acting system that connected the internal state of an actor to a formalized set of gestures and movements. This model was developed based on observations of human interactions across a range of situations [24]. The result of Marsella et al.'s work showed considerable consistency in the subjects' interpretation of given hand movements in animation based on Delsarte's rules.

¹ Laban movement notation is a system for understanding, observing, describing and notating all forms of movement for dance.

In addition to this work, Brenda Harger proposed an early study of using improvisational theatre models to develop believable characters. Specifically, she showed a simple animation of characters entering a room, where users can vary the characters' projected movements through one quantitative parameter: status. Through this parameter one can see different ways that characters can perform the entrance action [25, 26]. While Harger's work did not formalize a model for non-verbal behaviors, it built one step towards that goal by showing the effect of one parameter, *status*, on defining characters' posture, gaze, and mannerisms.

3 Our Study

We seek to extend the studies discussed in section 2.2 in search for a model that links non-verbal behavior to character attributes. To that end, we define three research questions:

1. Is there a set of character attributes, e.g., status, that are commonly understood by animators and can be used by designers to adequately describe a character?
2. Are there non-verbal behavior patterns that involve posture, gaze, and body movements with pacing and timing constraints?
3. Do variations in character attributes defined in 1 dictate distinct non-verbal behavior patterns defined in 2?

3.1 Character Attributes

What defines a character from a narrative and drama viewpoint? What attributes or parameters can be used to define such a character? These are still open questions. While there are several models available, they have not been computationally validated for the purposes of building believable characters. Previous research in interactive narrative used two models: Five Factor Model² [19] and a model based on traits [15]. In our view, the Five Factor model is very general and does not necessarily link well to non-verbal behaviors from a performance perspective. Character traits, on the other hand, are widely discussed, but have no standard definition.

Johnstone [27, 28] formulated two character models for describing a character for improvisational purposes; these models are: Fast-Food Laban based on the Laban movement notation [22] and Fast-Food Stanislavsky based on Stanislavsky's model [29, 30].

Stanislavsky is a famous Russian director who composed a theory for acting that is currently used by many acting schools to teach actors how to build and develop their characters. In his teachings, he discussed the importance of purpose for a character. Thus, instead of an actor playing an emotion, the actor would develop his actions depending on his character's goals, tactics, and purpose.

Johnstone then took the Stanislavsky model and developed several characters (examples are shown in table 1) defined in terms of purpose. He uses this model, which he named Fast Food Stanislavsky, in his improv exercises. In this paper, we propose to verify the utility of Fast-Food Stanislavsky regarding research question 1.

² A psychology-based personality model comprised of five personality dimensions: openness to experience, conscientiousness, extraversion, agreeableness, and neuroticism. This model was believed to explain all personality types [<http://www.personalityresearch.org/bigfive.html>].

Table 1. Two character definitions from Johnstone’s Fast-Food Stanislavsky

<i>To Give Someone a Bad Time</i>	<i>To be Thought a ‘Computer’</i>
<ul style="list-style-type: none"> • Invade their space. • Be restless, tap fingers. • Cross your legs away from them. • Frown; sigh; ‘tut’. • Glare at them. • Laugh at wrong time. • Poke them with finger. 	<ul style="list-style-type: none"> • Be cold and distant. • Be insensitive to pain or pleasure. • Dislike physical contact. • Other people are slow. • Pause before answering. • Be efficient – everything in its right place.

In addition, Johnstone identified several parameters of character representation that can affect non-verbal behaviors. One such parameter is *status*. *Status* is a major signifier that defines posture, gaze, and use of space in movement. A person of high status, for example, occupies more space, with erect posture, and always looks people in the eye. A person of low status, on the other hand, tends to occupy less space, with more inward posture, hunched back, and always looks away from people. Johnstone stated that characters often identify their status in comparison to other characters in the space, and behave relatively low or high by modifying their posture, gaze, and use of space in opposition to other characters in the space.

3.2 Study Design

In our study, we intend to verify the usefulness of the model defined above as well as identify a link between the model and non-verbal behavior patterns, addressing research questions 2 and 3. To this end, we recruited three animators. We gave them the task of animating ten variations of a simple two-character scenario, where the variations constituted variations in character definitions using Fast Food Stanislavsky.

While previous research targeting similar questions, such as [13], used actors to perform the scenario, we asked for the participation of animators rather than actors. We made this decision for various reasons. First, to gauge research question 1, we wanted to involve animators rather than actors, since the question involves the acceptance of the character model by animators. Second, while actors are good at performing through externalizing their internal feelings, animators are good at creatively thinking and composing of all facets of non-verbal behaviors. Third, actors think of only their part in relation to other characters. Animators, on the other hand, think of the scene as a whole and develop the characters in the scene from the ground up. We believe this is the first study we know of that was designed with participation of animators in the way discussed above. Thus, we will report on the process as well as the results in the discussion section.

3.2.1 The Scenario

The scenario is set up in an office environment with fixed dialogue between two male characters (see Table 2) [28]. We chose to use this scenario, because it has been successfully used by Johnstone [28] and Harger [26] to show variations in non-verbal behaviors.

Table 2. Simple Scenario taken from Johnstone's work [28]

Setting: Internal. An Office. Day.	
-	Officer (male): Come in, Smith. Sit down. I suppose you know why I've sent for you?
-	Smith (male): No, Sir. [Officer slides a newspaper over to Smith.]
-	Smith: I was hoping you wouldn't see that.
-	Officer: You know we can't employ anyone with a criminal record.
-	Smith: Won't you reconsider?
-	Officer: Good-bye, Smith.
-	Smith: I never wanted your bloody job anyway. [Exit.]

3.2.2 The Character Variations

The animators were given the scenario above and were told to animate it ten times with variations in characters using Fast-Food Stanislavsky (see Table 3). We chose to begin our study with only ten variations—a small sample, to validate the character model before we perform a full study. The animators were all given a 3D Maya file containing an office scene modelled with a desk, two chairs, a paper used as a prop for the scene, and two skeleton 3D characters models, which the animators used as their base for animation.

Table 3. Scene variations

Scene Number	Officer	Smith
1	High Status	Low Status
2	Give someone a bad time	To show people you are happy about everything
3	To show someone they are boring	To get sympathy
4	To flirt with someone	To accept guilt
5	To be thought normal	To flirt with someone
6	Low status	High status
7	To give someone a good time	To be thought a hero
8	To be thought a computer	To impress someone
9	To show people you are happy about everything	To show someone they are boring
10	To be thought intelligent	To be thought mysterious

3.2.3 The Animators

The animators recruited for this study are third to fourth year undergraduates studying at the School of Interactive Arts and Technology (SIAT) at Simon Fraser University.

They all completed the animation course required as part of the SIAT curriculum. The three animators were of different skill levels, namely professional, amateur, and beginner. While all animations produced for the project were of good quality, there were several differences in quality and assimilation of character descriptions that we attribute to the animators' varied skill levels. The professional animator works part-time at an animation company, and thus he was able to produce professional animations for the project. The amateur animator produced high quality animation, but had no industry experience, and thus his animations were not as good in quality as the professional animator. The beginner animator completed his animations, but it was obvious from our meeting notes that he was learning as he produced them. Videos of these animations will be presented at the conference.

3.2.4 The Process

The three animators worked independently on the initial Maya file to produce the scene variations described above. We had several group meetings: one at the beginning, one in the middle, and one at the end of their animation process. In these meetings, we clarified the confusions about the scenes and character variations; in the last meeting, we asked animators to share their experience and thoughts. These interactions were all documented as part of the study. The deliverables for each animator were ten scenes in Maya file format; we also asked them to produce ten rendered video clips of the animation with voice-over for demonstration.

We performed two kinds of analysis: high-level, namely, observations of animations, and low-level computational analysis of motion data. The former focuses on postures, actions, gaze, mannerisms, behavior habits, and character proximity, whereas the latter focuses on finding out details of timing, spatial relations, and movement of different body parts including head, arms, hands, and legs. In this paper, we discuss only the high-level analysis as the low-level analysis is still undergoing.

4 Results of the High-Level Qualitative Analysis

At the current stage, we have obtained some initial qualitative findings from the meeting notes and manual video coding of the animated scenes. The meeting notes informed us about the acceptance and appropriateness of the character models described by Johnstone (i.e. targeting research question 1) and the coding results showed some high-level patterns of character postures, gaze, mannerisms, gestures, unscripted actions as well as behavior habits as a function of the character model (i.e. targeting research question 2 and 3).

4.1 Validating the Character Model (Question 1)

Our first research question was designed to validate the character attribute model used. In part, we needed to verify if this model was understood by artists, specifically animators who will be involved in the design of interactive stories. This is an important step as it has implications on the use of this model as a tool for artists or designers to encode characters with improvisational ability within interactive stories.

Among the three versions (one for each animator) of the 10 scene variations, there were considerable consistency among the portrayal of specific characters, which indicates a coherent understanding of the character attributes of these characters and a well defined model as an indicator of non-verbal behavior. However, there were some minor inconsistencies. From our discussions and meeting notes, we deduced that the animators had difficulty portraying characters with such purposes as 'to be thought as hero', 'to impress someone', 'to be thought a computer', and 'to be thought mysterious'.

4.2 Non-verbal Behaviors (Question 2)

While the goal of the study is to identify patterns of non-verbal behaviors—sequence and parallel behaviors with time and spatial constraints, we report only on non-verbal behaviors here, since it is hard to quantify patterns and timing constraints qualitatively. In our next step, we will perform a computational analysis that will help identify non-verbal behavior patterns.

Among the ten scene variations, there were 15 character variations selected, among which 5 were used by both characters (see Table 3). During the video coding, for each scene variation we noted the postures, gestures, and actions that appeared in all three animators' works. In analyzing these variations and the meeting notes of all scenes, we found that 11 out of the 15 character models were consistently portrayed by the animators, i.e., these 11 models all have more than four noted consistent entries, be it posture, gesture, or action.

Among the 11 consistent models, 5 were those used by both characters. 3 out of these 5 character models were portrayed similarly for both Smith and the Officer. However, the other 2 character models showed different results in the animation. For the character model 'to show people you're happy about everything,' Smith was portrayed similarly by animators in one way, whereas the Officer was portrayed similarly by animators in another way. While both Smith and the Officer would gesture moderately and hold eye contact, they used different postures. Smith sat or stood, depending on the Officer's position. The Officer appeared to always touch the table and support part of his body weight on it. This pattern was convincing because the Officer was the owner of the space and naturally his posture showed ownership. Thus, we believe, in this case, that characters' power in the story context is a factor affecting non-verbal behavior. For the character model 'to flirt with someone,' however, our three animators portrayed the Officer quite similarly, but Smith very differently. In fact, our meeting notes showed that two animators had troubles imagining a flirting scene between characters of the same gender.

In opposition to the 11 consistent models, the rest 4 models, which were 'to be thought hero', 'to be thought a computer', 'to impress someone' and 'to be thought mysterious,' were less consistently portrayed by the three animators. The characters based on these models behaved either very differently, or in a rather inexpressive way. For example, there was a difference among three animators' portrayal of 'to be thought mysterious' character model. Two of them thought to be mysterious means not showing people a full self; therefore their character was either hiding behind objects or staying as far as possible. The third animator, however, considered a mysterious person a spy type; hence, the character showed curiosity and constantly peeked into the document the other character is reading.

Table 4 shows a segment of our results for research question 2. The table only shows non-verbal behaviors for two different character models that were consistently portrayed. The table shows non-verbal behaviors categorized in three dimensions: body motion in relation to self, body motion with interactions with props, body motion with interactions with other characters.³

Table 4. Results: non-verbal behavior as a function of character characteristics

Character	Motion Description
(Officer) To show people you're happy about everything	<p><u>General Body Motion:</u> [Posture] stand with fingers on table supporting some weight [Gesture] moderate amount; head moves when talking</p> <p><u>Motion with Props:</u> [Action] point and touch the paper when calling attention from Smith</p> <p><u>Motion in relation to others:</u> [Eye] hold eye contact most of the time</p>
(Smith) To show someone they're boring	<p><u>General Body Motion:</u> [Gesture] minimal amount with little actions (e.g. yawning, tapping, etc.)</p> <p><u>Motion with Props:</u> [Eye] follows when Officer is calling attention of the paper [Action] look at watch in the latter half of the conversation</p> <p><u>Motion in relation to others:</u> [Distance] far (table in between)</p>

4.3 Link of Character Attributes and Non-verbal Behavior (Question 3)

From the data table we obtained from the video coding, for each Fast-Food Stanislavsky character model we can conclude there were corresponding non-verbal behaviors used to portray this character. In other words, the non-verbal behaviors vary for each character model. We can reach this conclusion only for the models animated with consistency as described above. We believe the inconsistent models imply a varied interpretation, and thus cannot be used to deduce non-verbal behaviors.

5 Discussion

Our first step of data analysis was a manual coding of the video content of the 30 clips we collected from the three animators. From the above factual summary of the coding results, we are able to deduce *three* findings, which constitute the contribution of this paper. It should be noted that the study is still on going, and thus the findings are continuously growing as the study continues.

Our first finding focuses on a character model based on Johnstone's Fast-Food Stanislavsky (*research question 1*). 11 out of the 15 character models we tested showed a considerable degree of consistency among the three animators in terms of how they portrayed the non-verbal behaviors of the characters. This reflects an

³ The animation videos can be found at: <http://emiie.iat.sfu.ca/believablecharacters/videos/>

adequate degree of consistency of how animators understand and interpret these character models. Thus, we can assert that the Fast-Food Stanislavsky model can be used with some refinement as a character attribute model for interactive stories. With further studies, we can identify which attributes are consistently interpreted and which are not, and refine our model accordingly.

Using this model to indicate character instead of hand coding the animation presents two opportunities. First, it provides a faster content development cycle. Second, it provides improvisational space for virtual characters, i.e. it is a model with which characters can adapt their behaviors without reverting to hand coded routines. However, further experimentation with this model is required. Specifically, in the follow up study, we intend to examine ways of computationally encoding the model, its non-verbal behaviors, and algorithms for adequately firing and adapting the identified non-verbal behaviors to the context.

The second finding is concerned with the non-verbal behaviors identified (*research question 2*). Even though Johnstone listed many different non-verbal behaviors in his description of the Fast-Food Stanislavsky model, there are several details and non-verbal behaviors that were not fully discussed, specifically reactive and expressive actions, distance between characters, gaze, and how frequent the character gestures. These non-verbal behaviors are important in complementing the character models described by Johnstone. They are an extension from the existing acting rules and a stepping stone to defining the character non-verbal behavior patterns for computationally encoding character characteristics. As we continue with the computational low-level analysis of our animation data, we intend to develop patterns, which, as defined earlier, are lists of two or more movements with timing and spatial constraints.

The third and final finding is a list of two interesting lessons that we note from our qualitative analysis and meeting notes. First, while we treated each character separately in our discussion and study design, it was apparent that animators did not. They have indicated that they can show a character as intelligent for example by making the other characters in the scene impressed with what he/she is saying. Thus, this interaction between characters in the scene can also be a way of formulating a character in relation to others. Second, the recruited animators were all different in terms of their skill level, and thus their results also varied. Any model that we report on in the future will need to take these variations into account.

6 Conclusion and Future Work

The study described in this paper started with the premise that non-verbal behavior patterns can be identified as a function of character attributes. While there are many interesting findings noted in this paper, there is no theoretical model that can be concluded from this study. The road to such end requires several studies and exploratory experiments. The contribution of this paper is two-fold: (1) identify the problem and (2) present an approach to resolving the problem. The findings reported in the paper show success in the choice of the character attributes model and the beginning of the discovery of non-verbal behaviors that can be formulated as patterns linked to the character attributes defined.

Our job is still at its early stage. The next step is to computationally analyze the animation data, which will allow us to analyze the non-verbal behaviors in depth and identify low-level movement patterns related to each joint, which possibly can be grouped by body parts including head, eye, arm/hand and legs and which include timing and spatial constraints. To address the implications of the study, our future plans will involve more animators, ideally all professional ones coming from different education background, to ensure the generality of the data. We will also refine the description of each character variation, so that each item on the definition list is relevant to animating body motions, and change the wording when necessary.

Acknowledgements

Special thanks to our animators: Huan Chen, Harrison Wang, and Michael Chang for their animation works and permission to present their work. We would also like to thank David Milam, a graduate student at Simon Fraser University, who provided the animators with the initial Maya file of the office scene.

References

- [1] Rickel, J., Marsella, S., Gratch, J., Hill, R., Traum, D., Swartout, W.: Toward a New Generation of Virtual Humans for Interactive Experiences. *IEEE Intelligent Systems* 17 (2002)
- [2] Swartout, W., Hill, R.W.J., Gratch, J., Johnson, L.W., Kyriakakis, C., LaBore, C., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Morie, J.F., Rickel, J., Thiébaux, M., Tuch, L., Whitney, R.: Toward Holodeck: Integrating Graphics, Sounds, Character, and Story. In: *Proceedings of 5th International Conference on Autonomous Agents* (2001)
- [3] Seif El-Nasr, M., Saati, M., Milam, D., Neidenthal, S.: *Assassin's Creed - A multi-cultural Read. Loading* (forthcoming)
- [4] Cassell, J.: Embodied Conversational Agents: Representation and Intelligence in User Interfaces. *AI Magazine* 22 (2001)
- [5] Perlin, K.: Real-time Responsive Animation with Personality. *IEEE Transactions on Visualization and Computer Graphics* 1 (1995)
- [6] Perlin, K., Goldberg, A.: Improv: A system for Scripting Interactive Actors in Virtual Worlds. *Computer Graphics* 29 (1995)
- [7] Allbeck, J., Badler, N.: Representing and Parameterizing Behaviors. In: Prendinger, H., Ishizuka, M. (eds.) *Life-Like Characters: Tools, Affective Functions and Applications*. Springer, Heidelberg (2003)
- [8] Cassell, J.: A Framework for Gesture Generation and Interpretation. In: Cipolla, R., Pentland, A. (eds.) *Computer Vision in Human-Machine Interaction*, pp. 191–215. Cambridge University Press, New York (1998)
- [9] Aylett, R.S., Louchart, S., Dias, J., Paiva, A., Vala, M.: FearNot! - an experiment in emergent narrative. In: Panayiotopoulos, T., Gratch, J., Aylett, R.S., Ballin, D., Olivier, P., Rist, T. (eds.) *IWA 2005. LNCS (LNAI)*, vol. 3661, pp. 305–316. Springer, Heidelberg (2005)
- [10] Marsella, S., Gratch, J.: EMA: A Computational Model of Appraisal Dynamics. In: *Agent Construction and Emotions*, Vienna, Austria (2006)
- [11] Seif El-Nasr, M., Ioerger, T., Yen, J.: FLAME - Fuzzy Logic Adaptive Model of Emotions. In: *Autonomous Agents and Multi-Agent Systems*, vol. 3, pp. 219–257 (2000)

- [12] Ekman, P.: *Darwin and Facial expression: a Century of Research in Review*. New York Academic Press, New York City (1973)
- [13] Wallbott, H.G., Scherer, K.R.: Cues and Channels in emotion recognition. *Journal of Personality and Social Psychology* 51, 690–699 (1986)
- [14] Amaya, K., Bruderlin, A., Calvert, T.: Emotion From Motion. In: *Proceedings of the conference on Graphics interface* (1996)
- [15] Loyall, B., Bates, J.: *Personality Rich Believable Agents that Use Language*. In: *International Conference on Autonomous Agents* (1997)
- [16] Loyall, B.: *Believable Agents*. Computer Science Department, Carnegie Mellon University, Pittsburgh (1997)
- [17] Mateas, M., Stern, A.: A Behavior Language: Joint Action and Behavioral Idioms. In: Prendinger, H., Ishizuka, M. (eds.) *Life-like Characters. Tools, Affective Functions and Applications*. Springer, Heidelberg (2004)
- [18] Elliot, C., Rickel, J., Lester, J.: Lifelike Pedagogical Agents and Affective Computing: An Exploratory Synthesis. *Artificial Intelligence Today*, 195–211 (1999)
- [19] Andre, E., Martin, K., Gebhard, P., Allen, S., Rist, T.: Exploiting Models of Personality and Emotions to Control the Behavior of Animated Interactive Agents. In: *Agents 2000 Workshop* (2000)
- [20] Perlin, K.: Building Virtual Actors Who Can Really Act. In: Balet, O., Subsol, G., Torguet, P. (eds.) *ICVS 2003. LNCS, vol. 2897*, pp. 127–134. Springer, Heidelberg (2003)
- [21] Perlin, K.: Better acting in computer games: the use of procedural methods. *Computers and Graphics* 26 (2002)
- [22] Davies, E.: *Beyond Dance: Laban's Legacy of Movement Analysis*. Routledge (2006)
- [23] Marsella, S., Carnicke, S., Gratch, J., Okhmatovskaia, A., Rizzo, A.: An exploration of Delsarte's structural acting system. In: *6th International Conference on Intelligent Virtual Agents* (2006)
- [24] Zorn, J.W.: *The Essential Delsarte*. Scarecrow press, Metuchen (1968)
- [25] Harger, B.: Entertaining AI: Using Rules from Improvisational Acting to Create Unscripted Emotional Impact. In: *Game Developers Conference* (2008)
- [26] Harger, B.: Workshop on Improvisational Acting. In: *AAAI Symposium on Intelligent Narrative Technologies* (2007)
- [27] Johnstone, K.: *Impro: Improvisation and the Theatre*. Theatre Art Books (1987)
- [28] Johnstone, K.: *Impro for Storytellers*. Theatre Arts Books (1999)
- [29] Stanislavski: *An Actor Prepares*. Theatre Arts Books, New York (1936)
- [30] Stanislavski: *Building a Character*. Theatre Arts Books, New York (1949)

Revisiting Character-Based Affective Storytelling under a Narrative BDI Framework

Federico Peinado^{1,*}, Marc Cavazza², and David Pizzi²

¹ Facultad de Informática, Universidad Complutense de Madrid, Spain
email@federicopeinado.com

² School of Computing, University of Teesside, United Kingdom
{m.o.cavazza,d.pizzi}@tees.ac.uk

Abstract. Belief-Desire-Intention (BDI) is a well-known cognitive theory, especially in the field of Software Agents. Modelling characters using software agents has been proven to be a suitable approach for obtaining emergent and autonomous behaviours in Interactive Storytelling. In this paper it is claimed that an effective extension of previous models to the BDI framework is useful for designing intelligent characters. An example shows how internal thoughts and motivations of *Madame Bovary's* main characters can be more naturally formalised as a cognitive side of the story. A narrative reformulation of BDI theory is needed to avoid the implicit complexity of other proposals.

Keywords: Story Generation, Software Agents, Artificial Intelligence.

1 Introduction

Character-based approaches to Interactive Storytelling can be inspired on many different theories. A distinction which is particularly interesting in our research field is made between cognitive-oriented and narrative-oriented models.

While some researchers focus on simulations of human behaviour in which Storytelling is only a promising field of application, others are convinced that an excessive effort on modelling the rationale behind characters distracts us from the real problems of narrative phenomena. As revealed by current studies, a trade-off between narrative and cognitive orientations to character-based storytelling is needed in order to build a system capable of complex and deep storytelling with a balanced integration of these aspects. Belief-Desire-Intention (BDI) theory is starting to be applied in the field as a promising framework for that integration.

In this paper we study the use of BDI for modelling intelligent characters and develop our own *Narrative BDI* extension to an existing character-based storytelling system, in order to determine its real potential as an explanatory framework for the motivations of characters' behaviour.

* Funded by University of Teesside under the RIF initiative, Ministerio de Educación y Ciencia (TIN2006-14433-C02-01), Universidad Complutense de Madrid and Dirección General de Universidades e Investigación de la Comunidad de Madrid (CCG07-UCM/TIC 2803).

2 Related Work

The original Bratman’s theory of BDI [1] is a psychological model of human practical reasoning, created as an explanation of future-directed intentions in terms of current knowledge and goals of a rational agent. Different BDI architectures [9, 10] have evolved from this theory which bestows three mental attitudes for determining agent’s responses to perceived changes in the environment: beliefs, desires and intentions. *Beliefs* are facts about the world, including the agent itself, which may not necessarily be true. *Desires* are world states that the agent would like to accomplish, normally described as a set of goals. *Intentions* are internal states about what the agent is doing for satisfying its desires, usually executing or deliberating plans.

As the original theory suggests and following works explore [5] there is a complementary relationship between BDI and planning. Nowadays many storytelling systems rely on an automatic planner, implementing characters’ desires as *goals* and characters’ intentions as *plans*, which are composed by *operators* representing the sequence of actions that the character will try to perform following its desires.

2.1 Social Behaviour over BDI

Chang and Soo [3] claim that planning about how to influence the minds of others is a key piece to the generation of stories with sophisticated characters, as Bremond’s [2] pointed out before. For modelling characters with social behaviour they propose a BDI architecture in combination with an external “social planner”. Their agent-based story generator is based on Shakespeare’s *Othello*, which, as it is shown in Figure 1, contains a remarkable example of a manipulative character, Iago.

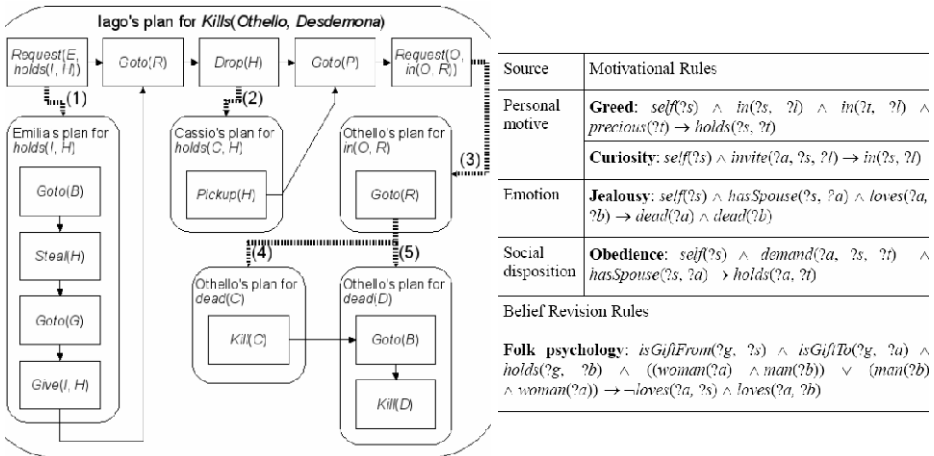


Fig. 1. Sample of the social planning proposal by Chang and Soo [3]

Social plans involve special actions of other agents apart from those than can be performed by their author. These special actions are called *foreign actions* as opposite to the *native actions* of the performing agent. Social plans do not use any social-specific

operators but usual ones as those corresponding to physical actions, e.g. go to Cassio's residence, or communicative acts, e.g. request Emilia to steal Desdemona's handkerchief. What are social-specific in this planning are the behaviour rules of Iago, the *social agent*. Some of these rules are used for *updating agent's beliefs about others* according to its perceptions, e.g. if Othello finds a gift that he has given to Desdemona in the Cassio's residence, he will infer she is loving Cassio, not him. Other rules are use for *updating agent's goals according to its beliefs about others*, e.g. if Othello thinks Desdemona is loving Cassio, he will desire to kill them both.

Foreign actions need their performer to be motivated in some way by the author of the social plan. *Preconditions* of each foreign action must include the requirement, , that some of its effects match some goals of the performing agent. By-passing the question of agents' perception, *effects* of both foreign and native actions can cause changes in other characters' beliefs and goals.

2.2 BDI, Emotions and Drama

Damiano and Pizzo [4] pointed out the importance of a mature cognitive model of characters' emotions as the ground on top of which the plot is developed. They distinguish between "character as agent" at the cognitive *level of actions* and "character as drama function" at the narrative *level of plot*.

In their work, a theoretical view of the BDI framework is enriched with that distinction and the use of emotions, being used to analyse a scene of Shakespeare's *Hamlet*. As shown in Figure 2, the protagonist is who really drives the scene, while Ophelia simply reacts to his actions.

<p>1.2 Hamlet at(Ophelia, nunnery) <i>Hamlet.hope</i></p> <p>Ophelia know(Ophelia, "Hamlet's feelings") - TRUE</p>	<p>1.2.1 Hamlet believe(Ophelia, "moral values are false") <i>Hamlet.anticipation</i></p>	<p>6 Hamlet introduce_topic(Hamlet,Ophelia,"moral values") - TRUE</p>	<p>ask (Requesting)</p>
		<p>7 Hamlet introduce_topic(Hamlet,Ophelia,"moral values") - TRUE</p>	<p>ask (Requesting)</p>
		<p>8 Hamlet declare(Hamlet,Ophelia,"moral values are false") - TRUE</p>	<p>declare (Stating)</p>
	<p>1.2.2 Hamlet believe(Hamlet,Ophelia, "affections are false") <i>Hamlet.expectation</i></p> <p>Ophelia know(Ophelia, "Hamlet's feelings") - TRUE</p>	<p>9 Hamlet argue(Hamlet,Ophelia,"moral values are false") - TRUE</p>	<p>do_nothing (EducationalProcess)</p>
		<p>10 Hamlet introduce_topic(Hamlet,Ophelia,"affections") - TRUE</p>	<p>declare (Stating)</p>
		<p>11 Hamlet declare(Hamlet,Ophelia,"affections are false") - TRUE</p>	<p>do_nothing (Stating)</p>
		<p>12 Hamlet argue(Hamlet,Ophelia,"affections are false") - TRUE</p>	<p>argue (EducationalProcess)</p>
		<p>13 Hamlet advise(Hamlet,Ophelia,(go(Ophelia,nunnery))) - TRUE</p>	<p>advise (Directing)</p>
		<p>14 Hamlet argue(Hamlet,Ophelia,(go(Ophelia,nunnery))) - TRUE</p>	<p>argue (EducationalProcess)</p>
	<p>1.2.3 Hamlet intend(Ophelia,go(Ophelia, nunnery)) <i>Hamlet.looking_forward</i></p>	<p>15 Hamlet advise(Hamlet,Ophelia,(go(Ophelia,nunnery))) - TRUE</p>	<p>do_nothing (Directing)</p>
		<p>16 Hamlet argue(Hamlet,Ophelia,(go(Ophelia,nunnery))) - TRUE</p>	<p>argue (EducationalProcess)</p>

Fig. 2. Sample of the BDI with emotions proposal by Damiano and Pizzo [4]

The methodology of these authors has two parts, a bottom-up analysis of the script and a subsequent reformulation of the story according to the proposed model.

The first step of their analysis is to identify the verbal and non-verbal *actions* at present time; then, *drama units* along the different layers of the plot tree should be recognised. Drama units have a *character* as subject and a *dramatic goal* associated to it which is connected to the frustration or the achievement (i.e. boolean result) of the corresponding agent's *physical goal*.

Their reformulation starts with the assignment of *emotions* to the agents, according to their cognitive states, after that, actions are mapped to *processes* and finally reorganised in a sequence of *beats*, i.e. action-reaction pairs. Characters are represented as a selection of those beliefs and emotions of their cognitive avatars that are relevant in the context of their dramatic units, i.e. high-level affective conflicts.

3 The Narrative BDI Model

This research is a continuation on previous work on Character-Based Affective Storytelling, on top of which our Narrative BDI Model is applied to enrich it with a cognitive layer. The referent is the Interactive Storytelling prototype published by Pizzi et al. [8], in which part of the novel *Madame Bovary* [6] is formalised.

The approach is based on a real-time planner whose operators mainly deal with *literary feelings* (as described by Flaubert himself [7]), instead of low-level mental or physical states, so they are called *emotional operators*. These are classified in three semantic categories. *Interpretation operators* update characters' feelings to respond to a change in the state of the world. *Interaction operators* intentionally modify other characters' beliefs. Finally *physical operators* basically change characters' locations.

The Narrative BDI model is a high-level narrative-oriented extension to the BDI approaches described in Section 2. As in previous approaches, only the relevant characters of a story require formalization as real BDI agents. Beliefs and feelings are directly modified by the effects of perceived emotional operators, the atomic actions of our representation. Desires and goals are usually updated in terms of character's feelings, the atomic facts. Intentions are divided in those that are part of a rational *plan* and those that are merely *reactions*. Both are composed only of interaction and physical operators because interpretation operators cannot be triggered by characters. These operators are under the control of the planner, representing that feelings are not usually updated by voluntary acts but according to characters' predefined personality.

The new formalisation of the *Madame Bovary* sample in terms of the Narrative BDI Model is summarized in Figure 3. Emma experiences the larger repertory of feelings and consequently is the protagonist in most of the operators, but Rodolphe is clearly leading these chapters of the novel from a rational point of view.

In *Madame Bovary* feelings play a dominant role and characters as Emma present modelling difficulties from the BDI point of view. While Rodolphe is an exemplary instance of a rational agent with pertinent beliefs about women's world, a clear goal and a detailed plan to achieve it, Emma lives a naive fantasy (even believing Rodolphe's lies), her desires are changeable and her behaviour is quite opportunistic.

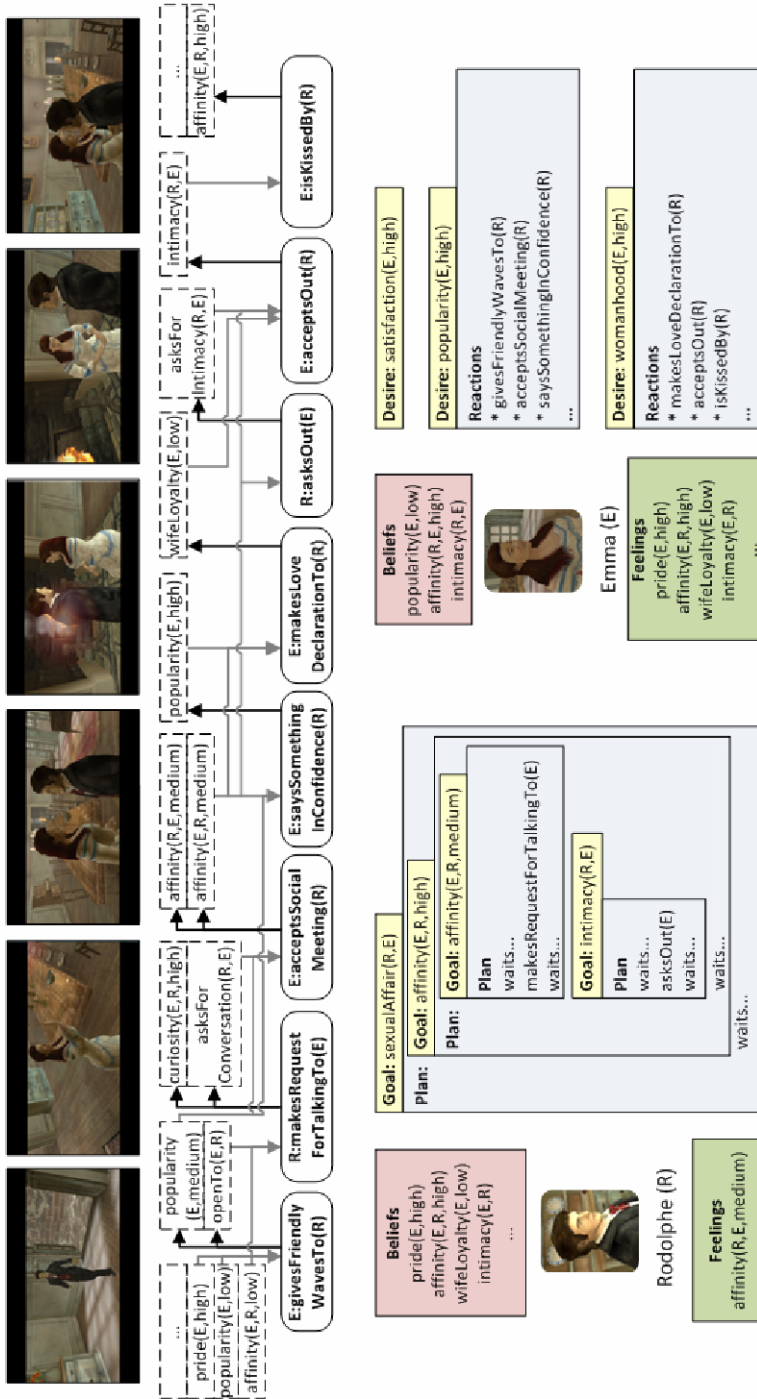


Fig. 3. Mapping the emotional planning approach to the Narrative BDI model

4 Conclusions

BDI theory is starting to be considered a promising tool for modelling sophisticated characters in Interactive Storytelling. During this theoretical analysis of the framework and its potential benefits for reformalising our character-based storytelling prototypes with motivated characters, four main advantages have been identified.

First, BDI as a cognitive model reinforces narrative causality insofar as motivations and beliefs are causal links that enrich characters. Second, the autonomy of software agents applies a more scalable structure to the way characters' beliefs and feelings are updated: BDI revision rules are distributed among the agents, which is not the case with planning operators. Third, the Narrative BDI model supports characters' behaviour visual explanation, illustrating in a more plausible fashion what is implicit and centralized in the sequence of emotional operators generated by our re-used planner. Fourth, although not explored in showed samples, multi-agent systems provide a natural approach to user interaction by simply adding a user-controlled agent to the repertory.

Most of the benefits derived from Chang and Soo's proposal (stronger social interactions and nested multi-character plots) are consequences of the interaction between many social agents, but as the reader should have notice neither of the three samples gathered here explore stories with more than one "full BDI" agent. Interoperability issues between the functionality provided by BDI platforms and external planners are also an interesting topic for further research.

References

1. Bratman, M.E.: *Intention, Plans, and Practical Reason*. CSLI Publications (1999)
2. Bremond, C.: *Logique du Recit*. Editions du Seuil, Paris (1973) (in French)
3. Chang, H.-M., Soo, V.-W.: *Planning to Influence Other Characters in Agent-Based Narratives*. In: *Integrating Technologies for Interactive Stories Workshop, International Conference on Intelligent Technologies for Interactive Entertainment*, pp. 12–17 (2008)
4. Damiano, R., Pizzo, A.: *Emotions in Drama Characters and Virtual Agents*. In: *Emotion, Personality, and Social Behavior*. AAAI Spring Symposium Series. Stanford University, California (2008)
5. de Silva, L., Padgham, L.: *Planning as Needed in BDI Systems, extended version of Planning on Demand in BDI Systems*. In: *International Conference on Automated Planning and Scheduling*. Monterey, California (2005)
6. Flaubert, G.: *Madame Bovary*. *La revue de Paris* (ed.), France (1856) (in French)
7. Leclerc, Y. (ed.): *Plans et Scenarios de Madame Bovary*. CNRS, France (1995) (in French)
8. Pizzi, D., Charles, F., Lugin, J.-L., Cavazza, M.: *Interactive Storytelling with Literary Feelings*. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) *ACII 2007*. LNCS, vol. 4738, pp. 630–641. Springer, Heidelberg (2007)
9. Rao, A., Georgeff, M.: *Modeling Rational Agents within a BDI-Architecture*. In: *2nd International Conference on Principles of Knowledge Representation and Reasoning*, pp. 473–484. Morgan Kaufmann, San Mateo (1991)
10. Saadi, A., Sahnoun, Z.: *Towards Intentional Agents to Manipulate Belief, Desire and Commitment Degrees*. In: *ACS/IEEE International Conference on Computer Systems and Applications*, pp. 515–520. IEEE Digital Library, Los Alamitos (2006)

***VirtualActor*: Endowing Virtual Characters with a Repertoire for Acting**

Ido A. Iurgel

CCG - Centro de Computação Gráfica and Centro Algoritmi, Universidade do Minho,
Campus de Azurém, 4800-058 Guimarães, Portugal
idoiurgel@yahoo.de

Abstract. The research project *VirtualActor* will develop Virtual Actors, with focus on creating linear 3D stories. Virtual Actors are virtual characters that can follow abstract directions on how to act, in analogy to a real actor that follows commands of a movie director. The Virtual Actors will be part of a framework called *Creator*, which will enable Virtual Actors that employ different AI technologies to be used together, thus providing a test-bed for virtual actor technologies. The authoring process will encompass innovative features that support an iterative and intuitive production of CG-movies. An important technology to be tested for virtual actor behavior control will be Case-Based Reasoning, which shall allow exploring the concept of a “repertoire” of behaviors for a virtual actor.

Keywords: Virtual Actor, Virtual Character, Character Animation, Authoring.

1 Introduction

Virtual Actors, in the sense of our research, are virtual characters that are employed for the creation of linear stories, i.e. of computer graphic (CG) movies. They are able to follow abstract commands of a human director that guide them to perform in a certain way. Virtual Actors could mean a considerable step forward for the production of 3D movies, because they can lower the production costs and allow everybody to express personal stories, without any specific animation training. Additionally, the technology of the Virtual Actors will contribute to novel kinds of video games and similar systems, including systems with specific interactive storytelling elements.

Up to now, Virtual Actors for movie production with specific expressive and acting faculties have not found ground in existing CG-production processes and applications. The project assumes that virtual actors for linear stories meet a clear demand for simplified animation tools. This demand arises with the emerging self-expression culture, where people put their photos, stories, and movies into the Internet. Many ordinary people are producing and publishing self-made art. Virtual actors as core elements of an application that facilitate the casual creation of CG-movies should be very welcome, in this situation. Virtual actors can also be employed to simplify and improve the animation process of commercial CG-movies, for example if the intelligent modules encode specialized knowledge on the interplay of gaze in conversation that facilitates the work or goes beyond the skills of an animator.

The research project *VirtualActor* will produce a framework for Virtual Actors that aims at enabling the creation and comparative study of Virtual Actors' technologies, enhancing their faculties and thus fostering their commercial and innovative use.

An early and important example of previous research is [1]. Recent examples (e.g. [2]) tend to give more weight to emotional models, and to focus on interactive narrative. Existing commercial applications such as Antics¹ enable the easy production of Machinima-like movies, without detailed AI-based control of the performance of the virtual characters. In contrast to these examples, *VirtualActor* focuses on the authoring process and on a generic framework for different technologies, and aims at the study of Virtual Actors with advanced communication and expressive faculties tailored for the iterative and intuitive production of movies.

2 The Project *VirtualActor*

This project starts in summer 2008², and lasts two years. It will follow this approach:

- **Primacy of authoring** and of the authoring process
- **A flexible runtime and authoring framework** that allows Virtual Actors with different technologies to interact, called *Creator* (Figure 1)
- **Multiple levels of direction**, with emphasis on the more abstract levels
- Possibility of **enhancing and correcting** the performance of the Virtual Actors **iteratively**.
- **A dedicated sub-system for the training phase** of the Virtual Actors

We will study faculties of a system that is centered on an iterative and intuitive creation process. Such a system should present, among others, the following properties:

- **Intelligent handling of underspecified situations:** Virtual Actors should never “do nothing”, but should instead improvise sensibly, employing available contextual information, during parts of the play that do not contain more specific directions. Thus, the author will always see a sensible performance, as a possible starting point for further enhancements.
- **Handling of dependence chains of behaviors:** A change in an aspect of the performance leads to the adaptation of a whole chain of related behaviors, e.g. a direction to actor A to look at another actor B, shall cause actor B to look back autonomously.

The first technology to be examined as AI-basis will be Case-Based Reasoning. The assumption is that the “cases” of a database of complex behaviors will enable the creation of a “repertoire” of surface acting for a specific Virtual Actor, and will provide the basis for the impression of this specific actor's personality and acting style. Thus, the training phase of such an actor would consist of associating *typical* acting patterns to *typical* story situations; the patterns will then be adapted and reused in similar situations.

¹ <http://www.antics3d.com/>

² The project is funded by the Portuguese research foundation FCT, Fundação para a Ciência e a Tecnologia, under the reference number PTDC/EIA/69236/2006.

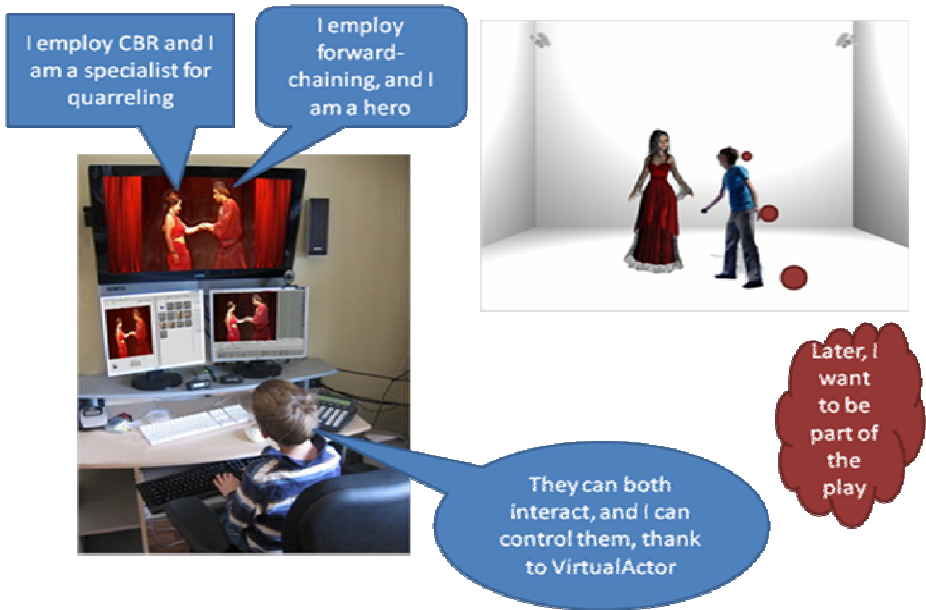


Fig. 1. The software framework of the *VirtualActor* project, *Creator*, will allow for interoperability of Virtual Actors with different underlying technologies

3 Conclusion

The project also assumes that understanding linear story production with Virtual Actors will lead to an understanding of the development of other types of real time reactive virtual characters. The Virtual Actor can also serve as a starting point for understanding the requirements and the functioning of an authoring tool that is suitable for other kinds of personality-rich intelligent virtual characters, i.e. also for the creation of actor-avatars, improvising actors, personality-rich assistants, and other kinds of interactive virtual agents (cf. [3]). This is because other scenarios of employing virtual characters will require the faculties of the Virtual Actors to express emotions, desires, and social relations, while taking into account higher-level commands and prohibitions.

References

1. Perlin, K., Goldberg, A.: Improv: A System for Scripting Interactive Actors in Virtual Worlds. *Computer Graphics* 29(3), 205–216 (1996)
2. Louchart, S., Aylett, R.: From synthetic characters to virtual actors. In: *Proceedings, 3rd Artificial Intelligence and Interactive Digital Entertainment International Conference (AIIDE 2007)*, pp. 88–91 (2007)
3. Iurgel, I., Marcos, A.: Employing personality-rich virtual persons - New tools required. *Computers & Graphics* 31(6), 827–836 (2007)

Steps towards a Generic Interface between Interactive Storytelling Applications and Character Animation Engines

Sebastian A. Weiß, Florian Berger, Alexander Marbach, and Wolfgang Müller

PH Weingarten, University of Education,
Media Education and Visualization Group (MEVis)
Leibnizstr. 3, 88250 Weingarten, Germany
{weiss,berger,marbach,mueller}@md-phw.de

Abstract. We present first results for the development of a general interface to believable agents. Central element of our approach is *CharanisML*, a character animation system meta language, which draws from a number of previous proposals for the description of interactive 3D content. To demonstrate the general character of this interface, we implemented two different clients in 2D and 3D that are able to interpret *CharanisML*. These clients may be adapted as animation engines for interactive digital storytelling engines like *Scenejo*. Using *CharanisML* it is possible for an author to control characters independently from both, storytelling engines and two or three-dimensional representations.

Keywords: Embodied conversational agents, interactive digital storytelling, behavioral animation, animation language, CharanisML, 2D, 3D, Scenejo.

1 Introduction

Telling a story requires the selection of an appropriate medium, taking in consideration the target audience and the context. In case of interactive digital storytelling (IDS) the main medium is computer-based. Still, there remain various alternatives for representing such stories, leaving choices in the use of 2D/3D perspectives, interaction, temporal behavior, and so on.

In this paper we focus on stories represented by the means of embodied agents. Still, different forms of representation may be required:

- **Web:** Mainly 2D
- **Online games and communities:** Second Life¹ or any other online game engine, mainly 3D
- **Desktop:** High quality rendering of 2D and 3D avatars

In practice, the decision for a representation also leads to a selection of a corresponding animation engine. However, due to this, previously authored story content has to

¹ <http://secondlife.com/> (last visited 04.07.2008)

be adapted when switching from one representation and corresponding engine to another. Preferably, it should be possible to author character control in general and not for a specific engine and to have the possibility to re-use already authored content.

There have been several approaches based on direct control of characters, or on simulation of native behavior concerning emotional and knowledge state during a speech act. In some situations an author wants to control the behavior of a character, in some not. We present a combined approach supporting both, direct control and autonomous animation based on predefined animation elements. We introduce two different client engines for rendering full bodied and model independent characters in 2D and 3D. These clients are proving our concept of a generic interface between storytelling platforms (in this case *Scenejo* [1]) and different animation clients.

2 Related Work

For character control several language standards and standard propositions do exist, like *FACS* for facial expressions [2] or *H-Anim* for abstract representation for modeling three dimensional human figures [3].

The Rich Representation Language (*RRL*) [4] developed in the *NECA* Project intends the fully automatically creation of animation based on words to be spoken. No animation control is possible without the text context.

The Virtual Human Markup Language (*VHML*) [5] combines several sub languages dealing with emotions, gestures, speech, facial animation, body animation, dialog data, and *XHTML*. However, specifications are fragmentary and there seems to be no further development.

AML [6] allows the definition of specific animation timestamps. It also provides means to synchronize specific actions. Yet, controlling character emotions is not directly supported; instead facial expressions are triggered directly. The speed of animations can only be set to slow, normal, and fast. With *AML* only avatars can be controlled. Scene, camera, and light source can neither be set, nor influenced.

3 Interface between IDS Applications and Animation Engines

We present our approach to a generic interface between interactive storytelling frameworks and animation engines. We called it *CharanisML*, which stands for **C**haracter **A**nimation **S**ystem **M**arkup **L**anguage. *CharanisML* is an XML-based command language to control remotely characters and scene environments. It relates to concepts from earlier works like *VHML*, *RRL*, and *AML*, and it partially supersedes them.

Analogue to the described paradigms, the *CharanisML* concept is based on three different levels of control:

- **System Control:** All kinds of system commands; parameter settings, etc.
- **Scene Control:** Adding of characters, objects, control of light source, and camera
- **Character Control:** Animation control, playing pre-defined animations, direct control

Its basic concept is that the storytelling application wraps a set of scene and character commands in a package and sends them to the character animation engine.

4 CharanisML Compatible Clients

In this section we present two implementations of *CharanisML* compatible animation clients. To prove the application of the language to both 2D and 3D, we implemented a client for two-dimensional comic style characters and a three-dimensional client supporting 3D scenes and characters.

CharanisML 3D Client

Based on the free *Ogre3D* rendering engine, the 3D Client supports *CharanisML* completely. For modeling characters, any modeling tool that is able to export to the Ogre file format can be used. Models of human figures have to be based on an H-Anim conform skeleton to be controlled.




Fig. 1. Left: Detail of 3D model with underlying skeleton. Right: Complete body model.

CharanisML 2D Client

The CharanisML 2D Client interprets the *CharanisML* core using a comic-like two-dimensional visualization. It is a cross-platform client and utilizes the *Shard* framework for interactive 2D environments. The 2D approach has the advantage of making character creation from scratch much easier compared to a 3D client. Using clipart or other templates, a satisfying character can be drawn comparably fast, without the need to care for skeletons or texturing issues involved with 3D design.

5 Summary and Conclusions

We presented first results on the development of a general interface to believable agents. *CharanisML* is based on a number of approaches for the description of 3D content. It is applicable for the use with 2D and 3D avatars. *Ogre3D* has been used as a game engine to build the 3D client. The cross platform media framework *SDL* is the base for the implementation of the 2D client. These successful implementations exemplify prove of concept of our proposition. Using *CharanisML*, an author may create character actions independently from an animation system. This means, it is possible to reuse content modules in another project requiring a different representation.

Future Work. Since the process of implementing the clients is ongoing, there are more features to be integrated. Currently, three-dimensional control of bones is ignored in the 2D client. A kind of interpretation layer for generating a two-dimensional equivalent has to be added. Nevertheless, control of 3D models may not be really adaptable to a 2D client. Nuances could not be visible or representable at all.

Acknowledgements. We want to thank Holger Biehl for the conceptual ideas behind and the development of the *Charanis* 3D client.

References

1. Weiß, S.A., Müller, W., Spierling, U., Steimle, F.: Scenejo – An Interactive Storytelling Platform. In: Subsol, G. (ed.) ICVS-VirtStory 2005. LNCS, vol. 3805, pp. 77–80. Springer, Heidelberg (2005)
2. Ekman, P., Friesen, W.: Facial Action Coding System: A Technique for the Measurement of Facial Movement. Consulting Psychologists Press (1978)
3. ISO/IEC 19774 — Information technology — Computer graphics and image processing — Humanoid animation (H-Anim) (last visited 04.07.2008), http://www.h-anim.org/Specifications/H-Anim200x/ISO_IEC_FCD_19774/
4. Piwek, P., Krenn, B., Schröder, M., Grice, M., Baumann, S., Pirker, H.: RRL: A Rich Representation Language for the Description of Agent Behaviour in NECA. Bologna, Italy (2002)
5. Gustavsson, C., Strindlund, L., Wiknertz, E., Beard, S., Huynh, Q., Marriott, A., Stallo, J.: VHML. Curtin University of Technology, working draft v0.3 edn. (2001)
6. Kshirsagar, S., Magnenat-Thalmann, N., Guye-Vuillème, A., Thalmann, D., Kamyab, K., Mamdani, E.: Avatar Markup Language. In: EGVE 2002: Proceedings of the workshop on virtual environments 2002, pp. 169–177. Eurographics Association, Aire-la-Ville (2002)

Looking at the Interactive Narrative Experience through the Eyes of the Participants

David Milam, Magy Seif El-Nasr, and Ron Wakkary

School of Interactive Arts and Technology,
Simon Fraser University
Surrey, BC, Canada
{dma35, magy, rwakkary}@sfu.ca

Abstract. The topic of interactive narrative has been under research for many years. While there has been much research exploring the development of new algorithms that enable and enhance interactive narratives, there has been little research focusing on the question of how players understand and internalize their interactive narrative experiences. This paper addresses this problem through conducting a phenomenological study on participants playing *Façade*; we specifically chose a phenomenological methodology due to its emphasis on the participants' lived experience from the participants' viewpoint. We chose *Façade*, because it is the only accessible example of an experience that revolves around social relationships, conflict, and drama as its core mechanics. In this paper, we discuss sixteen themes that resulted from the analysis of the data gathered through the study. In addition, we reflect on these themes discussing their relationship to participants' backgrounds, and project implications on the design of future interactive narratives.

Keywords: Design, Interactive Narrative, Storytelling.

1 Introduction

The topic of interactive narrative has been under debate for several years. What does it mean to be engulfed in an interactive narrative? Can users engage in a meaningful interactive narrative experience? Who tells the story, the designer or the player? While answers to these questions have not been formulated, the community is split.

A reasonable approach to this dilemma is to explore these questions through the design, development, and evaluation of interactive narrative experiences. Many researchers have explored the design of interactive narratives integrating believable agents [1], drama managers [2], user modeling [3, 4, 5], and planning systems [6]. In our view, the design of a good interactive narrative requires the understanding of the participants and their experience. Even though there has been much research on the development of interactive narratives, there has been very little research exploring how users view their interactive narrative experience. This paper focuses on a research study that attempts to understand the interactive narrative experience through the voices of the participants themselves, using a phenomenological method.

For the study, we chose to use *Façade* as an interactive narrative experience; *Façade* was developed by Mateas and Stern and released to the public in 2005 [7]. While some may argue that video and computer games are rich with examples of interactive narrative, we believe *Façade* is a better choice to explore. Most video and computer games use puzzles, quests, destruction, or collection as their core mechanics, where narrative is often used for motivation or game aesthetics. *Façade* focuses on social relationships, conflict, and drama as its core mechanics.

In this paper, we report results from a qualitative study exploring the question: what is the participant's experience of interactive narrative after playing *Façade*? Husserl's phenomenological philosophy [10] best suit our research question since it seeks a descriptive analysis of a phenomenon voiced by several individuals. Phenomenological methods do not seek statistical significance and instead emphasizes the participant voice in order to find a pure essence. For our purposes, we attempt to suspend judgment in order to articulate the essence of interactive narrative from the participants' in depth perceptions of their interactions. We chose to use a phenomenological method of data analysis to interpret the participant experience based on the works of Moustakas [8] and Colaizzi [9]. For this method a sample size of 7-12 participants is adequate.

The primary contribution of this work is in presenting results exploring the experience of *Façade*. In this paper, we describe the study we conducted, which resulted in sixteen themes describing the participants' experiences and thoughts after playing *Façade*. In addition, we will reflect on the relationship between these themes and participants' background, previous play experiences, and culture as well as discuss implications for future interactive narrative designs.

2 Previous Research

During the past few years there has been much research exploring the design of interactive narratives. Researchers within the Artificial Intelligence field, for example, have actively been seeking the development of new algorithms and systems that enable narrative adaptation and evolution as a function of users' interaction. The Oz project presented an early work in this area, where researchers concentrated on developing an interactive drama architecture composed of believable agents [1] with emotional responses [6, 11], and a drama manager that guides the drama as it unfolds [2]. Following their work, Mateas and Stern developed ABL (A Behavior Language), which allows designers to author character behaviors with joint goals. This language was used to author behaviors for *Façade* [7, 12]. In addition, few researchers recognized the utility of user modeling on drama management as a facilitator of conflict and drama [5, 13, 14, 15]. There are several research projects that explored different types of interactive narrative experiences, such as emergent narrative and third-person interaction models [16, 17].

The research works discussed above have a strong design and computational focus. Very few researchers focused on empirically evaluating the interactive narrative experience. From these few empirical studies, there are some who adopted a quantitative method evaluating their interactive narrative experience through likert scale questionnaires gauging specific areas of interest [16, 15]. Quantitative methods have several disadvantages, however, including constraining participants' responses to the questions

posed. Alternatively, other researchers explored using qualitative analysis to understand participants' experiences. Mallon and Webb [18] discuss a focus group-based research study that explores the users' interactive narrative experience within commercial role playing games. Their findings are interesting but pertain to a comparative analysis of four commercially available role playing games.

Two research projects have previously evaluated the experience of *Façade*. One study focused on evaluating the conversation interaction identifying participants' interpretation of conversation breakdown and character responses. They used qualitative analysis based on grounded theory where they triangulated data in the form of: observation notes, participants' interpretations of their actions after showing them the video of their interaction with *Façade*, and system tracing revealing the systems' inner interpretations of participants' utterances [19]. Some of their findings were similar to what we found in our study, as discussed later. The second study focused on evaluating participants' experiences across three different versions of *Façade*: two virtual desktop versions: in one version users type in their utterances and the in the other users speak their dialog; the third version is an Augmented Reality version where Trip and Grace are projected into the participants' physical space through an HMD (Head Mounted Display). They used qualitative analysis based on grounded theory to gauge the participants' sense of presence vs. engagement. Their results indicate that even though participants were more present in the AR *Façade* they were not as engaged as within the virtual desktop interface [20]. While these studies are closer to what we are exploring here, there are several differences. First, we present a phenomenology study of the users' experience of an interactive narrative. Thus, while we touch on many aspects of conversation (as discussed later), this is only one of the many elements we examine. Second, our subject pool is very different due to a different geographic location, culture, and school philosophy. Third, the study procedure and design presented is purely phenomenological in nature.

3 Façade

The story of *Façade* [21] introduces the player as a long time friend of Trip and Grace, two Non-Player characters, who have invited the player for an evening get together at their apartment. The participant takes on a first person perspective and interacts with Trip and Grace through natural language (as shown in Figure 1). The



Fig. 1. Screenshot of Trip and Grace with the participant interaction

player is also free to move about the apartment, manipulate objects, and perform simple gestures, such as kiss or hug. The game begins with him in the hallway outside Trip and Grace's apartment, where they can be overheard arguing. Once inside the apartment, he gets caught between Trip and Grace's arguments, as the drama unfolds. It is up to the player to resolve the course of the drama.

4 Study Design

4.1 Participants

We recruited eleven participants from the School of Interactive Arts and Technology (SIAT) at Simon Fraser University (SFU). All participants signed a consent form before they began the study. Since we didn't want to influence the study, we asked for specific participants who never played *Façade* before, but who are already familiar with new media and digital games. All eleven participants were undergraduates at SIAT enrolled in the Foundations of Game Design course. The diverse genres of games they enjoyed playing included action, role-playing, first person shooters, MMORPG, mobile, and early interactive fiction. Three of the participants have taken a course called Narrative in New Media. The average age was 24 years old; there were four females and seven males. The cultural background included six Canadians, two Canadian with Chinese decent, one Canadian with Japanese decent, one Iranian, and one Turkish. All participants received extra credit for their contribution in the study.

4.2 Procedure

The study included four phases. All phases occurred in one-on-one sessions taking place in a computer lab, lasting approximately one-hour. Phase I initiated within a one-on-one interview between the researcher and participants that took approximately 10 minutes. The interview started with an ice-breaker conversation, where the participant was asked to discuss his interests. He was then cued to talk about games he enjoyed playing, his cultural background, as well as his view of interactive narrative.

In phase II, the researcher asked the participant to read a description of *Façade* [www.interactivestory.net/#façade] and watch a YouTube clip [[youtube.com/watch?v=GmuLV9eMTkg](https://www.youtube.com/watch?v=GmuLV9eMTkg)]. The researcher then conducted a one-on-one interview exploring the question of interactive narrative again. This phase lasted around 10 minutes.

The researcher then proceeded to set up *Façade* for the participant to play for phase III. During this phase, the researcher refrained from speaking and instead took notes of noticeable actions the participant took and any other interesting observations (which were then noted as discussion points for phase IV). During play, participants were encouraged to speak about their experience, if they were comfortable to do so. Afterwards, we saved the recorded stageplay containing the session script of their interaction. This phase lasted around 20 minutes.

Phase four involved a *Façade* post-play interview. The interview questions were devised based on what the participants remembered from their play session along with their reflections. For example, the researcher asked participant "what stuck out in your memory", "elaborate on their non-verbal behavior such as laughing or

hand-tossing”, or “perhaps what they thought about their story ending?” These questions became catalysts to discuss other aspects of participants’ play experience. Due to complexity of data and their interdependencies, together with space limitations, we report on the results collected from the fourth phase only.

Many perspectives exist regarding validation of qualitative research data. In order for us to show validity, objectivity, and significance, we address the issue of credibility as discussed in previous research [22]. We address this issue first by including an external reviewer overseeing the analysis process and second, the analysis findings (in the form of theme descriptions) is sent back to participants to check for accuracy. To date, we have received positive responses from eight out of the eleven participants, agreeing with the analysis made and meanings formulated.

5 Analysis

We employed Colaizzi’s method [9] for analyzing the participants’ transcripts. Written transcripts were read several times to obtain an overall feeling for them. From each transcript, significant statements that pertain directly to their lived experience were identified. Meaning units were formulated from significant statements. The formulated meanings were then clustered into themes. We found 231 significant statements (each statement averaged 3 sentences long); we extracted these statements from over approximately 580 statements of the transcribed interviews conducted for phase four. Table 1 shows example significant statements (participants’ statements) and their formulated meanings (interpreted meaning). Arranging the formulated meanings into clusters resulted in 16 themes summarized in Table 2 and described below.

Table 1. Selected Examples of Significant Statements

Significant Statement	Formulated Meaning
I just wanted him to shut up and so I could talk to her, and force her to shut up so I could talk to him. But I wanted to get more in depth than they were allowing me to just because, I guess, the programming is looking for certain words.	Felt loss of control when interactions have little or no effect upon the story.
...in real life you kind of know your friends personality already. But in here it’s...although they are set to be “my friends”, I don’t know their personalities up until I interact with them.	Knowing a personality is drawn from previous lived experience. Learning about new personalities can take time separated from the course of the story.

Theme 1: Interactive Narrative is Not a Game. The conundrum of interactive narrative as a term has been a topic of debate. This was a frustration experienced by our participants as well, as they tried to define their own experience and relate it to the games they played. Five of the eleven participants made a distinction between *Façade* and games. Two discussed *Façade* as a “new form of entertainment” or “a story with game attributes.” One said it resembled the “real life situation.” Two found playing *Façade* a puzzle unfolding; one stated “puzzle[ing] it out,” while the other found “pok[ing] holes” in the system fun.

Table 2. Summary of 16 themes comprised from clusters of formulated meanings

Theme Descriptions – Phase Four	
1.	Interactive Narrative is Not a Game
2.	Interactive Fiction: Reading & Conversation
3.	Story Priming and Misalignment
4.	Story & Story Interaction
5.	Back-Story
6.	Character Believability (Action, Language and Comprehension)
7.	Participant Performance & Participant Interaction
8.	Clear Goals (narrative vs. Puzzle)
9.	Unsure of Control (narrative vs. Puzzle)
10.	Loss of Control – No Ownership
11.	Previous Lived Experience
12.	Social Participation (seeking to disengage)
13.	On Awkwardness
14.	Cultural Influences
15.	Replay Thoughts
16.	Testing the Boundaries

Theme 2: Interactive Fiction: Reading & Conversation. Three participants discussed the topic of conversation styles. They evaluated their experience with *Façade* in reference to popular games they played, such as *Kings Quest* and *World of Warcraft*. One found typing commands was similar to *King’s Quest’s* system, another pointed out the similarity between *Façade’s* conversation model and *WOW* chat conversation. Due to slight system lag in displaying the text on screen, when playing *Façade*, that participant stopped playing, he said “if there’s something the matter with the way I chat, then I give up ... I can’t continue to play because that is my voice.”

Theme 3: Story Priming and Misalignment. There were some misaligned expectations considering the resulting experience compared to *Façade’s* website described experience. After playing *Façade*, three participants expected the story to be different from the one experienced. This is partially due to how *Façade* was presented on the online website. One participant didn’t see how going back (to an old college friend) could lead to “this story that you wouldn’t expect.” This participant wanted Trip to explain how she introduced him to Grace. Another participant mentioned being “bi-ased” to think *Façade* supported comedic conversation opportunities, and that his dramatic story experience “did not match.”

Theme 4: Story & Story Interaction. Nine participants evaluated *Façade’s* story and their interaction with it according to its plausibility, model of interaction, and story flow. Two participants had divergent reactions on the plausibility of the story. One found the story ending “shocking” and “possible in real life,” while the other felt the story was “hard to tell” because the characters were too argumentative.

Seven participants discussed *Façade’s* unique model of player interaction with unfolding story and thought of ways to improve it. One participant found *Façade’s* conversation-based interaction “great” and more interesting than the story itself. This participant changed her affinity frequently and was “especially confused at the last part,” when Grace asks, “is what you’ve said tonight supposed to add up somehow, to something?”

Three participants mentioned difficulty with this model of interaction, as it continuously asked them to split their attention between following the story and taking the time to type responses. One was “so in the moment trying to get what’s going on between the two” that he felt like he missed many interaction opportunities. Another said, “I wasn’t sure if I should talk or what was supposed to happen because it was like tension building so I’m thinking do I break it or do they break it themselves.”

Another three participants emphasized more “meaningful” and “productive” interaction, for example, when “they [Trip and Grace] would ask me a question and, well clearly, I’m going to interact” but this would only serve to “piss the other one off.” Another two participants thought the story could be more interactive if the character offered to share activities, such as painting a picture together, or re-arranging the furniture.

Two participants mentioned playing their role by following the natural flow of the story since their initial responses were ineffective in stopping or changing the attitude of the argument. One reverted to this decision after he was kicked out the first time, while the other felt more immersed when he “just accepted it.”

Theme 5: Back-story. Five participants stated that they wanted more back-story for the characters. They wanted to know more about the characters personalities and get the “inside story” from one “point of view” or another. Knowing this would have helped them “choose proper words,” and facilitated a “more of an immersive” one-on-one dialogue. Three participants were interested to know or learn more about their own back-story “whose friend I was,” and “what kind of friend am I to them? I don’t know how deep my relationship is to them?” Knowing this would have better defined social “boundaries” especially when one participant tried to piece together why he was kicked out of the apartment “...they first want me to be involved in the conversation, but now they don’t want me to?”

Theme 6: Character Believability (Action, Language and Comprehension). Four participants commented on the characters’ (Trip and Grace) performance which centered upon their believability. One found their acting was “pretty good,” while another found Trip’s character to be “God awful” and “completely whiney.” One said, “...they make you feel like you’re talking to a person,” but were really “not listening.” Two felt they were “not reacting as people really would in a conversation” or “not listening,” because they “didn’t need me and didn’t answer me back half the time.” One exclaimed “are you reading what I’m writing?” Another voiced her heedless effort: “I was like sit down, calm down, you know listen; you’re not listening, listen to me, can I ask you a question all that just to be, you know (laugh).” Another participant was expecting a “better” reaction after repeatedly kissing the characters, which got him kicked out.

Theme 7: Participant Performance & Participant Interaction. Four players commented on their strategies. One player acted with a purpose to “egg them on” because she “had things to say ... I had things to say to both of them”, “I could be all nice-nice”, or “I could work Trip a little bit.” All four, however, addressed limitations to perform such as “I just wanted to get in there [the conversation]” and “you realize you’re the 3rd party in the room.”

The four participants wanted more control to “start some topic”, “change the subject”, “lead the conversation”, or “alternate the argument into something else in order

to see if I can get it to the point that I want.” One player was determined to change the tone of the conversation into a more positive one by saying “smile and be happy” and talking about “bunnies” because “I was hoping they would clue into what bunnies were.”

Theme 8: Clear Goals (narrative vs. puzzle). Five participants addressed clear, discernable goals as a strategy for success in the unfolding narrative. These participants associated a certain “function” to their role in an effort to “figure out a strategy” and “solve a puzzle.” One participant equated clear goals to having a “tangible reward or punishment” and found the story not “overly engaging.” Without clear goals one participant felt “this puzzle doesn’t listen to me” and another didn’t know if it was good or bad that “I didn’t know exactly what I should be doing. ... You’re trying to get involved in it or step away from it and they keep either pushing or pulling independent of what’s going on and you don’t really know where you might go with it.” One participant related to another game, where “you know what your ability is, what kind of role you are in, and why are you playing this role.”

Theme 9: Unsure of Control (narrative vs. puzzle). Seven Participants were unclear of how to control the narrative, although they knew they were controlling some aspects of it. “I wasn’t sure”, “I’m kinda confused”, and “I wasn’t really doing a lot”, are example comments. Another participant commented, “I was just typing and I don’t know how exactly it worked, whether it will just hear what I said to one or the other or it just kind of analyze what I said and make something happen. Yeah, I just didn’t know.” For another participant it was unclear to what extent he would be able to “stop Grace” or “change the situation,” since pushing the comfort tool too many times would yield a negative character response, which made him feel like “I am stopping the narration of the game.” Lastly one participant was unclear what to do after picking up the wine bottle “the fact that you could pick it up makes you think you could do something with it.”

Theme 10: Loss of Control - No Ownership. Ten participants felt aspects of their interaction had little or no effect on the story due to system constraints. One felt “it wasn’t my story at all, and it was like I had no part in it. It wasn’t about me and it wasn’t about anything I would know.” Another felt frustrated as he couldn’t stop Grace from walking out the door (and hence a dissolution of marriage). Although this participant was asked not to speak by Grace, he continued typing anyways: “go after her” and attempted to use the gesture options to “hug, kiss, and comfort.” Three participants felt “lost”, “left behind” or struggled to “regain control” in the story. One thought his actions were “interrupting” the story and said “I haven’t done anything, I was just there.” One commented using text conversations was “like I have a weapon, but I don’t know how to use it.” One said “I wasn’t even part of the conversation anymore...but I don’t want to be bzzzzz, bzzzzz each time;” another said “I could not break this conversation if my life depended on it,” and another said, “I wanted to get more in depth then they were allowing me to just because, I guess, the programming is looking for certain words.”

Four participants focused on the conversation pacing, similar to results discussed in the previous study on Façade conversations [19]. One commented that the pace was “really fast” and that the story wouldn’t “stall for you ... because too many things happened while typing.” Three participants elaborated upon their experience in other

turn-based games where “if you stall the game stalls,” or “my action should trigger the next interaction.” Some commented that they didn’t have enough “space to say my things;” they were contently “being cut-off”, as it takes them time to type or they lost the opportunity due to pacing.

Theme 11: Previous Lived Experience. Five participants associated their own real life experience in their understanding of the narrative. Three participants discussed “already knowing” your friends’ personality prior to a similar argumentative experience. This is important as it guides the “choice of words.” In *Façade*, they instead felt like they “didn’t know” the characters and didn’t identify with them.

Theme 12: Social Participation (seeking to disengage). The dramatic climate of *Façade*’s social situation discouraged six participants from fully engaging in the narrative. One was “really sensitive about negative energy.” Three were not motivated in the story; they made comments, such as “why should I even care about fixing a relationship?”, “I just wanted to let them figure it out”, and “I’m going to remove myself from the equation” to let them “work it out,” which still caused a “disturbing emotional effect.” Two participants were disengaged enough to want to “give up” and “get out” of the situation. One succinctly stated “I just don’t care” while another said “I felt like, I don’t know, like a poor friend who doesn’t know anything who doesn’t know how to help because she doesn’t know.”

Theme 13: On Awkwardness. The topic of *Façade*’s awkward social situation was addressed by five participants. Participants described example awkward moments for them, including the phone call, “being trapped between arguments”, “two people yelling at each other”, and “bickering” which made them feel “confused”, like “I don’t want to be here”, and “I don’t see where you were going with this.” One participant wanted to leave as soon as it became awkward because “in real life, I probably will not let myself get into that situation.”

Theme 14: Cultural Influences. Playing in a social relationship within *Façade*’s narrative had cultural implications for two participants of Japanese and Chinese descent. Regarding politeness, one said “I don’t think I should go around touching things,” which limited her environmental and character interactions including comforting. This participant felt she was unable to “touch” Trip and Grace even though this was one of the interaction features. This participant also preferred to remain quiet (not interrupt), and wait for the conversation to naturally end. She wanted to make some hot tea with Grace in the kitchen as a means to privately speak with Grace, although this strategy was not understood. Similarly, another participant wanted to take off his shoes upon entering the apartment, and said afterwards “it sets a barrier to tell me what is not provided.”

Theme 15: Replay Thoughts. Four participants discussed replay in relationship to their first or second playing experience. One equated “high replay value” to “intriguing characters and social dynamics” as seen in other narrative forms (i.e. books and games) although in this instance he would not replay due to themes pertaining to the loss of control and the awkward social situation. Another was interested in replay in order to explore “different resolutions” although later became more cynical in testing the boundaries of the system. Two participants wanted to replay to pursue distinctly

happier endings, however, upon replay both were dissatisfied “they’d argue with almost anything” and “it got back to the same dialogue... and you’re like ah crap.”

Theme 16: Testing the Boundaries. Three participants explicitly discussed breaking the boundaries or “not playing by the rules” of the system. Two of which were cynical and did not care to adhere to the intended narrative experience, because one perceived it to be a “social experiment” and the other treated it as a “comedy.”

6 Reflections and Conclusions

This phenomenological analysis resulted in an exhaustive description of the player-narrative interaction listed in the themes above. In this section we aim to discuss some recurring patterns and explain such patterns in relation to participants’ previous experience. As noted in our previous study [23], players’ experience and comments rely heavily on their background, previous experiences, and mind set. In our view, many of the themes described above can be traced back to the participants’ background which may influence implications drawn on future designs of interactive narrative. In addition, due to the small sample of participants within this study, the analysis of the themes in relation to their background and experience will give the reader an understanding of how such patterns can be understood and generalized.

One of the most recurring patterns within the themes discussed above pertains to the participant’s multifaceted experience of control within the interactive narrative. Almost all participants touched on this aspect in the themes above. For example, many participants felt no ownership and loss of control, because they could not easily identify with their suggested role and experienced problems with pacing and natural language, which led them to conclude that characters were not listening to them. To many this was a new form of interactive ‘puzzle’ that they couldn’t map to their previous gaming experiences. Some have tried to map *Façade*’s play experience to other games, such as *King’s Quest* and *Princess Maker*. These mappings created false expectations of clear goals and a puzzle with some “positive outcome”, which caused the experience of loss of control to be more pronounced. As one participant said, it is like having a *weapon* that you cannot use. With some more feedback participants may be able to learn and develop coping strategies to handle this new form of interaction.

Related to the point discussed above, the first impression that participants got from the website and first few minutes interacting with *Façade* led them to believe that they are free to write or do anything at any point in time. However, there were inner constraints that participants soon realized, which led to an aversion reaction and loss of control, as it is widely known in psychology that impression formation plays an important role on judgment and perception [24]. This is specifically apparent in themes 4 and 7, where some of the participants commented on the lack of strategies to corner one character. Also, as discussed in the *Façade*’s study reported in [19], several participants commented on the pacing and interaction: when they should type, when they should listen, how fast they should type before the characters move on to the next beat. A few participants also discussed the loss of control due to not knowing what words would affect the interaction.

Another recurring pattern pertains to the methods by which players internalize characters, including their personality, backstory, character traits, feelings, emotions,

motivations, and goals. This particular pattern surfaced in several themes including 5, 6, and 11. Participants indicated how knowing characters' back-story could facilitate their performance through informed interaction. A side effect of being sufficiently informed relates to how participants negatively assessed characters as "not reactive" or "not listening." Many participants made an analogy between this situation and real-life similar situations with friends or family. They described several inconsistencies between their previous experiences with such situations and their experience in *Façade*. For example, one participant noted that in their real-life experience, they would know their friends and thus would know how to interact with them. Others said in real-life they would just avoid such friends. These previous experiences shape their understanding and their engagement with an experience such as *Façade*.

In addition, another interesting finding from the themes discussed above involves how participants evaluated the social situation which informed how they interacted and engaged with the experience. Many found conversing on the topic of a doomed relationship or being stuck in an awkward situation unappealing. For instance, lacking social appeal led some participants to test the boundaries of the system rather than genuinely interact with the story. Playing a social situation is almost non-existent in previous forms or interactive models. This, thus, has caused much confusion and left many players feeling awkward or culturally removed.

It is worth mentioning that participants from different cultural backgrounds are also susceptible to miss-assess the social situation. For examples, subtle queues for interaction were missed for one participant due to her inability to interrupt other characters as interruption is considered impolite in her culture. These are examples of cultural norms that were expected within the minds of the participants as part of the social interaction norms, but were not facilitated within *Façade*.

In conclusion, we have presented sixteen themes that resulted from our analysis of interviews conducted with eleven participants after they played the interactive narrative *Façade*. In this paper, we presented all the themes to show the experience of playing an interactive narrative using participant's voices. The transcriptions of the interviews as well as all analysis phases were member checked by the participants themselves as well as reviewed by an external reviewer to establish validity. In addition to the themes, we also reflected on our understanding of the themes in regards to participants' previous experiences to help explain their attitudes, mind-sets, and thus actions. The contribution of this study is in the data presented as well as the methods used. We hope that this data can be used to influence future interactive narratives' design.

References

- [1] Loyall, B.A.: Believable Agents: Building Interactive Personalities. Computer Science Department, PhD. dissertation, p. 222. Carnegie Mellon University, Pittsburg (1997)
- [2] Weyhrauch, P.: Guiding Interactive Drama. In: School of Computer Science, PhD. dissertation, Carnegie Mellon University, Pittsburg (1997)
- [3] Szilas, N.: Stepping into the Interactive Drama. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 14–25. Springer, Heidelberg (2004)
- [4] Seif El-Nasr, M.: Interaction, Narrative, and Drama Creating an Adaptive Interactive Narrative using Performance Arts Theories (2007)

- [5] Magerko, B.: Story representation and interactive drama. In: 1st Artificial Intelligence and Interactive Digital Entertainment Conference, Marina del Rey, CA (2005)
- [6] Bates, J., Loyall, B., Reilly, S.: An Architecture for Action, Emotion, and Social Behavior. *Computer Science*, 14 (1992)
- [7] Mateas, M., Stern, A.: Façade: An Experiment in Building a Fully-Realized Interactive Drama. In: *Game Developers Conference*, San Jose, CA (2003)
- [8] Moustakas, C.: *Phenomenological Research Methods*. Thousand Oaks, CA (1994)
- [9] Colaizzi, P.F.: *Psychological research as the phenomenologist views it*. Oxford University Press, New York (1978)
- [10] Husserl: *Ideas: General introduction to pure phenomenology*. Northwestern University Press, Evanston (1931)
- [11] Reilly, S.: *Believable Social and Emotional Agents*. Computer Science, PhD. dissertation, Carnegie Mellon University, Pittsburg (1996)
- [12] Mateas, M., Stearn, A.: *Procedural Authorship: A Case-Study Of the Interactive Drama Façade*. In: *Digital Arts and Culture (DAC)*, Copenhagen (2005)
- [13] Szilas, N.: *The Future of Interactive Drama*. In: 2005 Australasian conference on interactive entertainment (2005)
- [14] Seif El-Nasr, M.: *An Interactive Narrative Architecture based on Filmmaking Theory* (2004)
- [15] Thue, D., Bulitko, V., Spetch, M., Wasylishen, E.: *Learning Player Preferences to Inform Delayed Authoring*. In: *AAAI Fall Symposium on Intelligent Narrative Technologies*, Arlington, VA (2007)
- [16] Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M.: *FearNot! - An Experiment in Emergent Narrative*. In: Panayiotopoulos, T., Gratch, J., Aylett, R.S., Ballin, D., Olivier, P., Rist, T. (eds.) *IVA 2005. LNCS*, vol. 3661, pp. 305–316. Springer, Heidelberg (2005)
- [17] Pizzi, D., Cavazza, M.: *Affective Storytelling based on Characters' Feelings*, University of Teesside (2007)
- [18] Mallon, B., Webb, B.: *Stand up and take your place: identifying narrative elements in narrative adventure and role-play games*. *Computers in Entertainment (CIE)* 3, 6 (2005)
- [19] Mehta, M., Dow, S., Mateas, M., MacIntyre, B.: *Evaluating a Conversation-centered Interactive Drama*. In: *Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2007)* (2007)
- [20] Dow, S.M., Lausier, A., MacIntyre, B., Mateas, M.: *Initial Lessons from ARFaçade, An Interactive Augmented Reality Drama*. In: *ACM SIGCHI Conference on Advances in Computer Entertainment (ACE 2006)*, Los Angeles (2006)
- [21] Stern, M.M.A.: *Architecture, Authorial Idioms and Early Observations of the Interactive Drama Façade* (2002)
- [22] Lincoln, Y.S., Guba, E.G.: *Naturalistic Inquiry*. Beverly Hills, CA (1985)
- [23] Seif El-Nasr, M., Al-Saati, M., Niedenthal, S., Milam, D.: *Assassin's Creed: A Multicultural Read*. In: 2nd Annual Canadian Game Studies Association Conference, Vancouver, Canada (2008)
- [24] Pennington, D.: *Social Cognition (Routledge Modular Psychology)*. Routledge, London (2000)

Play and Narration as Patterns of Meaning Construction: Theoretical Foundation and Empirical Evaluation of the User Experience of Interactive Films

Regina Friess

GamesLab, Fachhochschule für Technik und Wirtschaft Berlin
Treskowallee 8, 10318 Berlin, Germany
Regina.friess@fhtw-berlin.de

Abstract. The central scope of the research project is a theoretical qualification (and empirical evaluation) of play and narration as different patterns of meaning construction. Based on cognitive film theory and theories of play an opposing specification of playful and narrative perception patterns is developed. Playful and narrative meaning construction are understood as interdependent constructs whose interferences/ interactions are to be analyzed. The theoretical concept is empirically surveyed by comparing the user experiences of a linear fictional film with those of two interactive versions of the same film. The initiating hypothesis is that interactivity enforces playful perception patterns and diminishes the relevance of narrative patterns for the users' experience.

Keywords: Play versus narration, playful versus narrative meaning-construction, user experience of interactive stories/ interactive films, meaning construction and user experience.

1 Introduction

The central scope of the research project is a theoretical qualification (and empirical evaluation) of play and narration as different patterns of meaning construction. Preliminary, notions of interactivity and play as linked concepts shaped the framework. The aim was to evaluate on an empirical basis if an interactive form produces playful reception patterns,¹ and how this effects on return existing narrative forms of perception.

The underlying assumption is that any kind of user experience is based on an active process of meaning making, which is connected to the concrete form of the media product.² The evaluation of the user experience in terms of pleasure or suspense should therefore be based on a preceding evaluation of meaning construction.

In a comparative analysis the core components of play and narration as forms of meaning construction are discussed and qualified on a theoretical basis. The two concepts are perceived as polar dimensions, which can intersect and correlate. It is

¹ These reflections refer to different theoretical approaches on interactivity and interactive media. See, e.g. [1].

² Based on theoretical approaches of symbolic interactionism and constructivism. See, e.g. [2], [3].

understood that there are possible dominations of either a more narrative or a more playful form of perception, which is to be specified.

The following research tasks were formulated:

1. How can playful and narrative patterns of reception be specified on a theoretical basis
 - a. in terms of meaning construction and
 - b. with regard to the derived qualities of involvement?
2. The empirical evaluation of meaning construction and involvement analysing the reception of a linear film compared to interactive narrative films based on these concepts.

2 Theoretical Framework

In order to analyse the process of meaning construction in dependence of the textual quality (interactive vs. linear) different levels and components of the reception process were defined. The separation of user activity and interpretation process and the interdependencies of the two levels is regarded as crucial concept in this framework. The presupposition is that the meaning construction (in cognitive and emotional respects) in an active context differs from that in a passive perception context. It is assumed that most of the existing research in interactive narration suffers from a pre-conditioned equivalence of meaning construction in linear and interactive stories as acting in a given fictional world and neglects possible interdependencies of user activity and meaning construction in the first step.³

The aim is to specify narrative and playful patterns and the correlating forms of interpretation and emotional involvement with regard to previously separated components of the perception process.

2.1 Specification of Play and Narration as Patterns of Meaning Construction and Derivation of Specific Forms of Involvement

The qualification of narrative perception patterns and components of interpretation and involvement is mainly based on cognitive and psychological film theories outlined by Edward Branigan, Ed Tan, Murray Smith and David Bordwell [7-10]. As there is hardly any theoretical work on play as perception pattern⁴, this part is based on theoretical reflections of different fields of play theory.⁵

³ The idea of the user acting in a given narrative context, without questioning the construction of the diegetic world as precedent process, is stated in several studies focusing on concepts of media effects in relation to given expectations and motivations (like e.g. Zillmans entertainment theory.) See e.g. the surveys of Vorder/Hartmann/Schramm [4] and Schlütz [5]. Also in Klimmt's level based concept of enjoyment of playing digital games the framing narrative level is discussed in terms of generalised narrative contents: "On this level, the players most often participate in a narration (...)" [6].

⁴ The research of Ohler/Nielsing [11] in playful patterns of film reception seems to be one of the rare exceptions.

⁵ Amongst others: the philosophical perspective of Klaus, Georg [12] and the psycho-sociological analysis of Sutton-Smith [13].

Table 1. Overview on play and narration as patterns of meaning construction and correlated forms of involvement⁶

Patterns of play in media perception	Patterns of narration in media perception
1. Interpretation	1. Interpretation
<ul style="list-style-type: none"> • domination of relational meaning construction • flexibility of external reference (operational meaning) • two-stage process of primary and secondary meaning construction (reframing) • textual coherence based on structural qualities • process of divergence • formal and pattern based linking 	<ul style="list-style-type: none"> • domination of referential meaning construction • explicit external reference (eidetic meaning) • semantical form of building reference focussing on human activity and events • textual coherence based on semantic reference to real life experiences • process of convergence • domination of causal linking of meaningful units
2. Involvement	2. Involvement
<ul style="list-style-type: none"> • self centred involvement in active meaning construction • focus on self experience • act of creating diversity and variety • joy of exploring abstract concepts, structural relations • gratification of feeling of competence 	<ul style="list-style-type: none"> • external focus on understanding a given meaning • focus on the experience of others • act of narrowing down the most possible meaning • interest in thematic fields, concrete actions and events • gratification of empathetic understanding/ sympathy
3. Distance/sovereignty	3. Distance/sovereignty
<ul style="list-style-type: none"> • distance by abstracted reference to reality • control of selfdirected activity 	<ul style="list-style-type: none"> • distance by reinterpretation of narrative point of view • control of cognitive and emotional dedication

3 Empirical Study

3.1 Central Hypotheses and Research Design of the Empirical Study

The initiatory hypothesis is that the interactive quality leads to an increased relevance of playful patterns in meaning construction and user experience. Based on theoretical

⁶ The specifications refer to idealized concepts, not to actual processes. The reference for the narrative pattern is the perception of a simple (nonfictional) kind of narration. The table gives a reduced overview on the main aspects without explicit separation of the considered levels of involvement (like general aspects of meaning construction, process-related aspects, the act of media usage, communicational aspects).

reflections on meaning construction and involvement during the perception of linear films, assumable changes correlated to an increase of playful reception patterns were derived.

The underlying understanding is that an alignment with the characters as possible real persons and the involvement in their actions and needs is a central concept for the user experience (of an all time drama movie). Based on the different characteristics of playful reception patterns (described above) the main suppositions are that the user's perception is more self centred, engaged in the process of constructing different possible versions (and therefore less empathetic).

The central hypotheses of the experimental study were:

- The interactive movies are perceived as less realistic and the user's awareness of the artistic quality is more distinctive.
- The characters are considered more functional, and less personal. Empathic involvement is diminishing.
- According to the degree in empathy, the involvement in the course of the story is lower.

The empirical survey has been carried out in summer 2006 at different universities in Berlin. There were three experimental groups ($n = 37, 38, 47$), each watching a different variation of a short fictional movie, one linear version and two slightly distinct interactive versions. The two interactive versions allowed the users to choose some actions of the protagonists. At four plot points, the film was frozen and the users could choose whether either the male or the female main protagonist should act (Fig.2). In the third version an additional score was shown evaluating the user's choice (Fig. 3 and Fig. 4). The plot and interactive structure were identical to version 2. The participants answered a standardised questionnaire and additional group interviews (with 7 to 9 participants) were performed for each version.



Fig. 1. Version 2 with choice



Fig. 2. Version 3 with scores after choice

3.2 Main Results from the Empirical Study

The evaluation of the standardised questionnaires indicated irritating and contradicting results in the first place. The second version achieved the best marks (see Table 1) for all main items (entertainment, realism, suspense, and empathy) whereas the linear

and the third version (with scores) polled notably worse and reached about the same ratings. Different control variables (age, gender, digital gaming experience, genre preferences) were checked and gave no results.

Table 2. Evaluation of the film versions: Means of the Main Items

Items (short titles)	Means (1 = very good5 = very bad)		
	Film 1 (linear) n = 47	Film 2 (interactive) n = 38	Film 3 (with scores) n = 37
Overall evaluation	3,0	2,5	2,9
Empathy/ Kirsten	3,4	3,2	3,6
Empathy/ Kai	2,8	2,4	2,8
Authenticity Characters	3,4	2,9	3,4
Realism	2,9	2,5	2,9
Suspense	3,3	2,8	3,3

Besides the unexpected high evaluation of the second version compared to the linear version, the more difficult task was to explain the differences between the second and the third version.

The analysis of the interviewees allowed for a further insight into the underlying thoughts and motives of the participants. The evaluation of the guided interviewees focused on two aspects: A clustering of the statements as it regards the content according to the evaluation of the main items, and an additional evaluation of the quality of the statements with regards to the degree of meta-reflection on the film construction or abstraction of an immediate discussion of plot events. As the total account of the interview participants is relatively low, the interpretations do not base on representational data, but can be seen as hints for further research.

In contrary to the previously made assumption the participants of the second group were notably deeper engaged in the characters' personalities, motives and actions than those of the linear version. The statements and discussions were mainly preoccupied with the characters and their activities. Differently the discussion in the first group was more directed towards a judgement of the film as a whole perceived from the perspective of a spectator. In the third version the preliminary demand to help the loving couple to come together and the scores seemed to initiate the predicted changes towards a more self-centred and functional perception. The participants were mostly preoccupied with reaching their goals and showed a more distant perception of the characters. The bad evaluation seemed to be caused by the combination of a lower involvement in the fictional story and an unfulfilled desire to manipulate the story according to own ideas.

4 Conclusion

Despite the unpredicted results of the questionnaire, the theoretical construction of polar dimensions of narrative and playful perception patterns are estimated as

consistent concepts. The analysis of the interviews seemed to underpin the hypothesis of changing meaning construction due to different alignments of the users' perspectives. The highly distinctive evaluations of the second and third versions based on very small changes demonstrate the relevance of specific design components for the user experience. This leads to the conclusion that further research in the relationship of interactive design qualities and meaning construction based on playful or narrative patterns seems promising.

References

1. Krämer, S.: Spielerische Interaktion: Überlegungen zu unserem Umgang mit Instrumenten. [Playful Interaction: Considerations on Our Usage of Instruments]. In: Rötzer, F. (ed.) *Schöne neue Welten? Auf dem Weg zu einer neuen Spielkultur*, pp. 225–236. Boer, München (1995)
2. Krotz, F.: Der symbolisch-interaktionistische Beitrag zur Untersuchung von Medien-nutzung und –rezeption [A Symbolic Interactionist Contribution to Research of Media Use and Media Perception]. In: Hasebrink, U., Krotz, F. (eds.) *Die Zuschauer als Fernsehregis-seure? Zum Verständnis individueller Nutzungs- und Rezeptionsmuster*, pp. 52–75. Nomos, Baden-Baden/ Hamburg (1996)
3. Schmidt, S.: Konstruktivismus als Medientheorie [Constructivist Media Theory]. In: Nöth, W., Wenz, K. (eds.) *Intervalle 2: Schriften zur Kulturforschung*, pp. 21–47. Kassel University Press, Kassel (1998)
4. Vorderer, P., Knobloch, S., Schramm, H.: Does Entertainment Suffer From Interactivity? The Impact of Watching an Interactive TV Movie on Viewers Experience of Entertainment. *Media Psychology* 3, 343–363, 344–345 (2001)
5. Schlütz, D.: Bildschirmspiele und ihre Faszination: Zuwendungsmotive, Gratifikationen und Erleben interaktiver Medienangebote [Videogames and their Fascination: Motivation, Gratificaton and Experience of Interactive Media], pp. 37–39. Munich (2002)
6. Klimmt, C.: Dimensions and Determinants of the Enjoyment of Playing Digital Games. In: Copier, M., Raessens, J. (eds.) *Level Up: Digital Games Reserach Conference*. Utrecht, pp. 246–257 (2003)
7. Branigan, E.: *Narrative Comprehension and Film*. Routledge, London (1998)
8. Tan, E.: *Emotion and the Structure of Narrative Film: Film as an Emotion Machine*. Lawrence Erlbaum, New Jersey (1996)
9. Smith, M.: *Engaging Characters: Fiction, Emotion, and the Cinema*. Clarendon Press, Oxford (1995)
10. Bordwell, D.: *Narration in the Fiction Film*. Routledge, London (1985)
11. Ohler, P., Niding, G.: Antizipation und Spieltätigkeit bei der Rezeption narrativer Filme [Anticipation and Play during the Reception of Narrative Films]. In: *Beiträge zur Film- und Fernsehwissenschaft*, vol. 60(42), pp. 13–30 (2001)
12. Klaus, G.: *Spieltheorie in philosophischer Sicht*. [Play Theory from a Philosophical Point of View] VEB, Berlin (1968)
13. Sutton-Smith, B.: *Die Dialektik des Spiels*. [The Ambiguity of Play]. Karl Hofmann, Schorndorf (1972)

Trying to Get Trapped in the Past – Exploring the Illusion of Presence in Virtual Drama

Georg Struck, Ralf Böse¹, and Ulrike Spierling²

¹ FH Schmalkalden, University of Applied Sciences, Blechhammer,
D-98574 Schmalkalden, Germany

² FH Erfurt, University of Applied Sciences, Altonaerstr. 25, 99085 Erfurt, Germany
georg.struck@gmail.com, r.boese@fh-sm.de,
spierling@fh-erfurt.de

Abstract. This article presents the prototype of a virtual cultural heritage installation that employs dramatic storytelling to give visitors a sense of personal presence in a reconstructed historical situation. A cinematic story is delivered through stereoscopic video and spatial audio featuring human actors in a virtual environment. This ‘light’ version of virtual reality offers a chance to study some mechanisms of immersion available to virtual story designers. The immersive effect of stereoscopic vision and 3D audio is explored and opportunities to heighten the effect through dramatic storytelling are discussed. Trading a real-time perspective for the presence of ‘holographic’ human actors, the installation focuses on the impact of dramatic narrative in an immersive environment. The article suggests that virtual storytelling allows designers to take the spectator’s illusion of personal concern with the proceedings one step further than traditional drama, enabling the spectator to feel personally present and thus ‘physically’ concerned. The opportunities and limitations of this spatially ‘situating’ effect are evaluated in the context of the installation.

Keywords: Immersion, presence, drama, virtual storytelling, virtual cultural heritage, stereoscopic video, wave field synthesis, spatial audio.

1 Introduction

The ruins of the monastery Georgenthal in southern Thuringia pose a special task for a virtual cultural heritage installation. Hardly visible, they consist mostly of basement walls and some artefacts that weren’t stolen during the looting and pillaging of the monastery in the 16th century. Still, the remains bear testament to one of Germany’s most important Cistercian monasteries founded in direct filiation to Morimund in France in 1140. The monastery was deserted in 1525 during the peasant wars and ultimately destroyed in 1550 (see Fig. 1 and 2).

‘Virtual Cultural Heritage’ applications are gaining popularity with the technical ability to create virtual worlds and represent historical knowledge in virtual storytelling. Digital reconstructions are particularly appropriate for simulating historical architecture that no longer exists. Up to now, virtual cultural heritage applications have been mostly employed in museums to convey factual historical knowledge.

The prototype presented here was tasked with introducing visitors to the history of the monastery, setting the stage for their exploration of the museum grounds. The team of InnoTP [4] wanted to steer clear of dry lecturing and tried a different path to convey knowledge about the monastery of Georgenthal.

- Universal human concerns with war, fear, peace and love were chosen as the core of the narrative, identifying suitable themes that would bridge the gap between present-day spectators and the protagonists. The goal was an emotional connection that would create a context of relevance for the historical subject matter.
- Taking the personal perspective of a novice monk in 1525, the narrative integrates knowledge about the architectural layout of the monastery, the routine of Cistercian monastic life and the disruption of monastic society in times of reformation and the peasant wars.
- The narrative attempts to maximise the emotional impact of the installation, using elements of suspense, mystery and romance and delivering the drama with intense stimuli to the visual, acoustic and haptic senses. The narrative relies particularly heavily on the spatial audio component, aiming to deliver the narrative also to audiences with visual impairments.

A simple observation gave rise to a special goal for the installation: the monastery is only visible to visitors through technological mediation. This creates an unusually equal situation for visitors with handicaps, particularly for visitors with visual impairments. The team decided to design the media in a fashion that would keep this basic equality, creating a common experience for all audiences. Thus, the immersive audio technology ‘IOSONO’ [5] was chosen as the main medium of delivery, charged with more than just complementing the visualization. IOSONO offers precise spatial hearing based on wave field synthesis [6] and the script was specifically written to leverage its potential.

The open character of the installation is echoed in the museum layout developed by architects Biessmann & Büttner for the ‘Stiftung Thüringer Schösser und Gärten’, inviting handicapped visitors into a ‘park of the senses’.



Fig. 1. Left: Hall under the abbot’s residence of monastery Georgenthal. Right: Hall under the abbot’s residence today.

2 Technical Platform

The technical platform combined stereoscopic computer graphics, live-action video and the spatial audio rendering system IOSONO, being the first to integrate these technologies.

2.1 Digital Scene Modeling

Based on the available architectural sources, the entire monastery premises were reconstructed with their architectural attributes as a 3D model at the Fachhochschule Schmalkalden's Department of Informatics. (Fig. 2 shows an interior view of the basilica created in Autodesk Maya.)




Fig. 2. Partial view of the virtual basilica

Conflicting appraisals of architectural sources and the remaining walls had to be resolved with the help of historians. The script's interior scenes called for additional research on the interior decoration and lighting of the monastery. An archive of historical views of the monastery and similar buildings was established and several artefacts were analysed on site to generate an accurate depiction of the buildings.

To avoid visual discrepancies between the computer-generated and live action scenes, the models were based on photographs taken on location. The photographs were later modified to be used as textures and to guide the generation of additional textures. To record virtual scenes stereoscopically, two virtual cameras were connected in Maya. The cameras were synchronized using MEL and enabled the camera man to easily change important stereo parameters, e.g. the distance of the eyes.

Large-screen projection and compositing with live-action footage made high resolutions necessary and the computational demands increased accordingly. The database with models, textures and project data had an overall volume of 100 GB managed by a Current Version System (CVS). With more than 2.5 Mio polygons in the full model, the computer-generated scenes had about 10.500 frames with 10 minutes render time each, totalling 1.750 hours for the final animation.

2.2 Integration with Live-Action Footage

The script called for real actors to move through historical settings in the monastery. Alternative locations could be found for close-ups; but if the actors were to be seen against the virtually reconstructed backdrop, the action had to be shot against bluescreen, keyed and fused with matching camera takes of the virtual background.

The majority of the film was shot in a blue box at Fachhochschule Schmalkalden. If possible, medium and close-up shots were recorded on location in Georgenthal and in the *Liebfrauenkirche*, a contemporary church in neighbouring Arnstadt.

Two identical camera models were mounted on a tripod using a special armature and controlled through a synchronization unit. But matching the white balance and aperture remained a challenge. Also, since a good deal of luminance is lost in stereoscopic projection due to the necessary filters, the scenes had to be lit with more light than usual, making it harder to maintain a consistent cinematic style. The peculiarities of stereoscopic camera work are summarized in [3].

When the footage reached the compositor, two camera streams needed to be matched, as well as two layers of virtual and real imagery. The perspectives of virtual and real camera had to be aligned, the stereoscopic camera settings used on each take (parallax angle etc.) adjusted and the different colour temperature readings of the cameras corrected. The script's many low contrast night-time scenes were a special challenge for compositing, often making it difficult to key blue screen scenes and fuse the virtual lighting with the real.

Editing was further complicated when it turned out that production with the spatial audio system required a complete lockdown on timing.

2.3 Integration with the Spatial Audio System

In wave field synthesis (WFS) systems, a circular array of loudspeakers is controlled through a spatial audio rendering application to create the illusion of sound sources located all around the listener. Virtual sound sources can be played back as point sources, focussed sources or planar waves. A point source emits sound in all directions and allows listeners to locate it exactly in space. Planar waves are experienced as coming from roughly the same direction and are mostly used to model acoustic spatial geometries. Focused sources are special – they are used to position virtual sound sources inside the circular loudspeaker array and will only be perceived correctly inside a specific area of sound reconstruction. Combining these sound sources, WFS systems create the illusion of discrete sound sources moving around and even through the audience.

A particularly impressive example in the installation is the attack on the monastery. Set in a phase of darkness, the visitors can locate each voice of the attackers and victims moving past and through each other in surprising acoustic transparency.

To prepare for a WFS production, all audio sources should be recorded as dry as possible in order to synthesize the necessary spatial acoustics independently. The WFS system then needs all the spatial geometries, including the coordinates of dynamic subjects. A direct export from the animation application is theoretically possible but was not yet available at the time of production [2]. The system then calculates the sound fields according to timecode. Any changes require a recalculation.

Special challenges arose from the combination with stereoscopic projection, as some spectators could end up with diverging acoustic and visual localization of objects. Special research addressing the issue concluded that the effects expected in the chosen application would be within the realm of tolerability [2].

2.4 Integration with the Stereo Projection System

Different systems for stereo projection have been developed [1]. The team chose the Infitec technology relying on passive projection encoded with interference filters and a pair of glasses for decoding [9]. In this set-up, the images are projected onto a silver screen with two projectors while filters help to separate the images for the right and left eye. These filters move the light colours on the spectrum towards the ultraviolet for one eye and towards the infrared for the other. The filter glasses decode the light appropriately for each eye.

One advantage of this projection technology is that the channels for each eye are kept in perfect separation. There are no ‘ghost images’ and every surface can be used for projection, ruling out expensive metal coatings. Also, spectators can move their heads freely without impoverishing the image quality. The disadvantages are that two projectors are necessary, the amount of projected light is reduced by the filters and the glasses are more costly than polarisation glasses.

The installation was realized as a prototype in the Fachhochschule Schmalkalden. The prototype offers seating space for 15 visitors surrounded by a circular array of 128 IOSONO WFS-loudspeakers and 4 additional bass-speakers, projecting from the ceiling onto a sound-transparent screen in complete darkness.

3 Story Design

The main story objective was to heighten the sense of personal presence in the monastery. The team wanted to avoid the transparent virtuality of a flythrough video and reach for a more complete suspension of disbelief. By experiencing a dramatic story in the virtual building and creating a personal emotional bond with the reconstructed monastery, visitors should feel like they had actually been there.

Theme. In a first step, the development team searched for a theme that could bridge the gap between past and present. The theme needed to be sufficiently universal to concern humans today and 500 years ago, being of existential emotional relevance.

The history of the monastery soon offered a starting point that could create an instinctive connection. Monastery Georgenthal was attacked and pillaged in the peasant war. Triggered by Martin Luther's revolutionary writings, the peasant wars initiated a period of religious fighting that would last for more than 150 years, ravaging large portions of Middle Europe. Sadly, fear of religiously motivated wars is also an emotion shared by many people today. Fear of war was thus chosen as a theme, since it elicits deep-rooted, instinctive emotions and concerns humans then as now.

Protagonists. The contemporary monk Georg Spalatin had already been discovered early on as one of the most interesting characters connected with the monastery. The former novice master had laid the foundation for his impressive career assembling a unique library of antique manuscripts in the monastery. Later when the monastery was attacked, he was already an influential supporter of Luther, a translator of important theological and philosophical books and a key figure of German Humanism.

On closer inspection, it turned out that Spalatin offered a productive link to the chosen theme. He had sensed a mounting threat of war and tried to counteract with a translation of the 'Querela Pacis' by Erasmus of Rotterdam. In this essay, literally translated 'lament of peace', Erasmus personifies peace as a godly woman talking to mankind, endowing her with an enigmatic female personality and beautiful voice, at once gentle, wise and sad. She reminds humans of their privileged, divine nature and warns them not to create unnecessary suffering for a shadowlike, apparitional world.

The goddess of peace develops a neo-platonic view of the world, rendering absurd all possible motivations for war like greed, anger, revenge or fear: "On the door to your very palace, at the gates of your fortress, behold the threshold to eternity! Why then do you fight over shadows, ghosts, apparitions of a waking dream, as if this life were endless, and time enough to be angry or unhappy."

Plot. Based on the chosen theme, a plot was developed with a second *ingénue* protagonist living through the process of insight demanded in the Querela Pacis. The young monk Domenicus is caught in the attack on the monastery in 1525 and seeks a way out of the threatening wartime situation. The threat is aggravated when he is imprisoned in the refectory as an alleged thief and cannot flee with the other monks. Defenceless, he hears the aggressive peasants overrun the monastery walls and enter the premises. In this passage, the story elicits fear and suspense in the audience.

Looking for an escape, the protagonist meets the mysterious figure of Georg Spalatin. Meanwhile a chaplain at the Wittenberg court, he is here to save the monastery library from the ransacking peasants. Charismatic Spalatin is fearless in an irritating way - the threat seems meaningless to him.

He frees the young monk from prison and offers him an escape: his student Magdalena will save him from the besieged monastery if he manages to meet her in the library. But his help is tied to a mission. Spalatin hands him a newly illustrated copy of the 'Querela Pacis' and asks him to take it to safety while he talks to the peasants. Though the protagonist is initially desperate to escape, the meeting with Spalatin and his beautiful student establishes a new goal, opening up a constructive perspective: instead of escaping, the young monk can now actively serve the purpose of peace.

Character Arc. This shift in goals sets the stage for the protagonist's character arc and transformation. The American screenwriting professor Frank Daniels describes the growth process of protagonists in drama by distinguishing between an initial, rather simple goal (want), in this case the protagonist's urge to escape; and a second goal (need) that emerges as result of a learning process, in this case the calm service of a higher purpose [8]. The protagonist finds a different kind of refuge in the 'Querela Pacis'.

Objectifying the new goal, Spalatin's edition of the 'Querela Pacis' makes the service of peace physically actionable in the story. Visually, the goddess of peace is further personified in Magdalena, Spalatin's young and extraordinarily beautiful student, lending a sensual dimension to the protagonist's new goal. The goddess of peace is voiced by the same actress, filling the room with a soothing, disembodied voice every time the book is opened ("Source, mother, patroness and protectress of life am I. Yet, I can find easier refuge with the wildest of beasts than with this one animal, the rational, immortal animal called man.")

Theme-related Twist. The protagonist abandons his blinding fear - and discovers that he was the victim of an illusion. The nightmare of being chased by attackers in the monastery was so frightening that he mistook it for reality. Offered a resort that is in keeping with Erasmus of Rotterdam's neo-platonic world-view, he realizes that he merely has to wake up from the 'the shadows of a waking dream'. In truth, he is safely in the present.

To enable this solution, a background story was introduced, presenting the young monk in a short scene at the beginning of the movie. He finds a stone fragment in the present-day ruins, close to the site of the installation in the old corn house. Returning to this starting point, the protagonist realizes that the stone fragment triggered his intense daydream when, again, the disembodied voice of the peace goddess is heard. It turns out that she made him relive this story to reveal the hiding place of the lost 'Querela Pacis' edition.

The young monk 'remembers' and finds the lost illustration in the corn house. Every stone tells a story, the visitor is reminded, or, as Georg Spalatin puts it in his closing words: his monastery is still contained in each stone and waits to be built anew. Without having to tell a time-travel story, this solution offered a chance to end the story in the present, creating a parallel between visitors and protagonist. A time-travel solution was abandoned early in the development process as it would have directed the attention at fantasy elements such as portals, tunnels etc.

The stone fragment that triggered the protagonist's daydream is present in the installation, lying below the screen on a table with old books. Connecting the fictional and the real in the sense of Murray's 'threshold object' [7], the stone offers a tactile dimension, prepares visitors for their own walk through the museum premises and centres the attention on the story's true protagonist, the monastery of Georghenthal.

4 Storytelling for Immersion – Exploiting the 'Situating' Effect

The story amplified the immersive medium's effect of feeling personally present and 'situated' in space by introducing a threat. The attention of the audience was instinctively focussed on their exact location in space as they found themselves stuck in a

dangerous place, threatened with violence. Playing with the archetypal nightmare of being caught in a house with robbers in the night, visitors were induced to constantly locate and monitor the threat – and the script made this task particularly difficult.

4.1 Leveraging the ‘Situating’ Effect of Stereoscopic Video

Different from film, the spectator in stereo projections is not an untouchable, godlike eye floating above the action. He inhabits a specific point in space. Compelled to track and evade the attackers in the installation, the visitor was in intense contact with the virtual space, scanning the building for hiding spaces and ducking from peasants. The script made use of this effect in two scenes where the protagonist hides in a niche, then jumps into the courtyard, pressing up against a wall and hiding just below the windows. Making the specifics of the location count, visitors could be seen ducking in the audience, trying to reduce their visibility.

Sharing a space with ‘holographically present’ humans strengthened the illusion that the world was ‘real’, offering a glimpse of virtual storytelling beyond the uncanny valley.

Like in a thriller, the attackers were hardly visible, prolonging the phase of suspenseful dread and letting the threat grow in the audience’s imagination. While the audience fled from hiding place to hiding place, the attackers were announced by other terrified monks running past, could be heard abusing other monks or guessed from the light of torches passing by.

The story played with the visitor’s urge to see – and denied them a free perspective. In one scene, the imprisoned protagonist listened in on Spalatin’s conversation with Magdalena. The protagonist’s perspective was reduced to a view from a narrow window – and his object of interest remained just outside the angle of view.

The story forced the audience to become aware of its own presence in space, repeatedly threatening and trapping them in this location.

4.2 Leveraging the ‘Situating’ Effect of Spatial 3D Audio

A unique feature of the IOSONO system was used to amplify the feeling of being personally exposed to a threatening environment. Other surround sound systems treat the entire seating area of an auditorium as one spot in fictional space, rendering all sounds for a single acoustic perspective. The IOSONO system offers no such protective bubble for visitors. Accurately placing the auditorium’s dimensions in fictional space, a sound source can enter the seating area and move between members of the audience. This effect comes as an unsettling surprise to the audience when the peasant’s break into the monastery – and the attackers also breach into the safety of the auditorium. No longer untouchable, the attackers seem to storm right through their midst, their precise localisation making them seem disturbingly real.

The script builds towards these moments, letting the attackers initially only brush past the auditorium, piercing the bubble in a few disorienting jabs and finally flooding the seating area with attackers.

Likewise, the story’s soothing forces are also introduced stepwise into the intimacy of the auditorium. In a key scene, mentor figure Spalatin suddenly seems only inches from the visitor’s ears (‘You have to free your mind of fear’), later he seems to circle

around the audience or talk from their midst. Building towards a counteracting effect, Spalatin and the peace goddess are rendered acoustically omnipresent in the climactic scenes, engulfing and relaxing the audience from all sides.

4.3 Reflecting on the Immersive Effect

The illusion of being personally present is reflected in the story itself. Amused, Spalatin warns the protagonist to take the dramatic events too seriously: the world around him is nothing but a ‘waking dream’. The visitor finds his own viewing situation brought up in the story – and can perceive the stereoscopic projection as perfect illustration of Spalatin’s neo-platonic world-view: though apparently real, spaces remain insubstantial and people move around the audience like ghosts.

Spalatin identifies fear as the main reason why the insubstantial is taken for real. He should rid his mind of fear, Spalatin warns the protagonist when he meets him for the first time, suddenly appearing in his prison cell. The image drops and it grows dark. When the image returns, Spalatin stands on the other side of the room. Distorted by fear, the visual sense can be misleading. From now on, whenever scenes grow too frightening, e.g. when the protagonist leaves his cell and steps out into the dangerous cloister, the image drops, leaving the spectator in darkness. After an initial increase in fear, the visitor realizes he is in the safety of the auditorium, surrounded by an instable world shaped by emotions rather than physics.

This ‘counter-dramatic’, disillusioning force reaches a climax towards the end when the peace goddess talks in complete darkness with her soothing, omnipresent voice, providing the calm needed to detach from the binding power of the nightmare.

Playing with the spectator’s suspension of disbelief, dropping the visitor in and out of the immersive world, the installation reflects upon its own mechanisms of immersion to illustrate the story’s historical theme.

5 Evaluation and Discussion

5.1 Lessons Learned

The ‘situating’ effect of the medium offers new opportunities but also imposes limitations on the director. The rough cut showed that short takes were incompatible with the effect, leaving too little time for the audience to ‘situate’ themselves. Worse, faster edits left visitors feeling restless and stressed. Viewers needed 1-2 seconds before their minds could establish the illusion of spatial depth and perspective. Any story information given during this initial phase was lost, creating considerable discontent and an inability to follow the story.

In our experience, stereo productions should avoid takes of less than 5 seconds and strive for long tracking shots of 20-30 seconds. Giving audiences the chance to savour the plasticity of the virtual space, long shots with a fore-, middle- and background seem to work best, while close-ups can feel disappointingly ‘flat’.

Static shots seem to have less impact. If the camera doesn’t move, at least the object of interest should. Particularly impressive are movements towards the audience.

5.2 Formative Evaluation

So far, about 200 participants visited the installation. Video observation, short questionnaires and in-depth interviews with 86 participants yielded first insights into the general effect of the prototype.

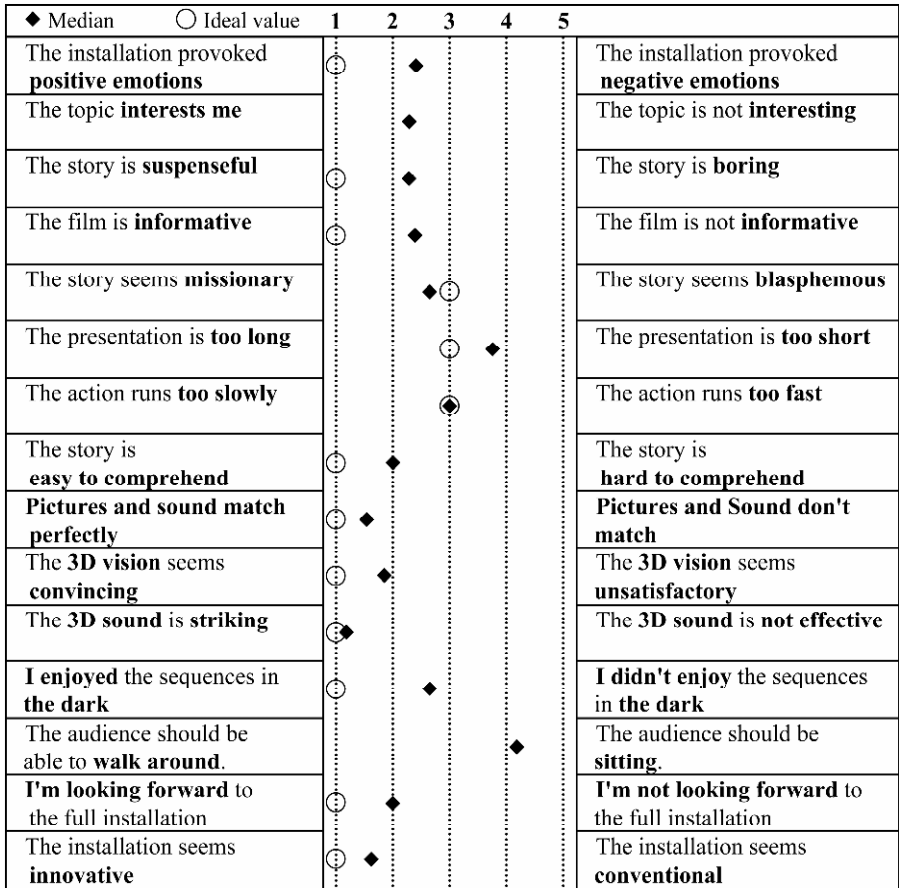


Fig. 3. Overview of evaluation results (a potential ideal value is noted as an empty circle)

Rather than in a statistical result, we were interested in individual and actionable criticism to improve the prototype. Therefore, in addition to questionnaires, we used focus group discussions. In particular, the effects of spatial audio, stereo vision and the emotional impact were examined.

While a detailed summary of this qualitative evaluation would exceed the limits of this article, the salient points were that the most visitors were affected and intrigued by the spatial audio feature. However, the stereoscopic perspective left many feeling strained and unable to focus. The evaluation also points towards dramaturgical flaws in the plot (too long, too complex) weakening the general acceptance of the installation. Further research into the experiences of handicapped audiences is in preparation.

6 Conclusion

In the presented Virtual Cultural Heritage application, immersive media were combined with a dramatic story to give audiences a strong sense of personal presence, creating a feeling of ‘having really been there’.

By choosing a theme that concerns present-day audiences, the story offers visitors a personal emotional experience in the virtually reconstructed monastery. It places the audience in the monastery only minutes before it is attacked in the 1525 peasant wars. Dealing with the threat of war, it triggers deep-rooted emotions and takes visitors on a personal quest for peace.

Dramatic storytelling amplifies the ‘situating’ effect of the stereoscopic projection and spatial audio by inserting visitors in a space that is invaded by violent attackers, creating special concern for their exact location in relation to the aggressive peasants.

The story personifies the monastery’s historical value in the characters of humanistic scholar Georg Spalatin and a figure from his favourite book, Erasmus of Rotterdam’s *lament of peace* ‘Querela Pacis’. Discovering an alternative to fearful escape, the audience takes part in the evacuation, loss and rediscovery of a precious book, shining a light on the monastery’s heritage of humanistic pacifism.

Acknowledgements

The project *InnoTP* has received funds from the German Research State Department BMBF. Thanks to all academic and corporate project partners, especially Biessmann & Büttner (architects), Fraunhofer IDMT (wave field synthesis), the students of the Fachhochschule Schmalkalden and the actors Ariana Antunovic (Magdalena), Alexander Döring (Thomas) und Uwe Zerbe (Spalatin).

References

1. Herbig, G. P.: Virtuelle Wahrnehmung der Tiefe bei der Stereoprojektion (Last access: 18.9.2008), <http://www.herbig-3d.de/german/tiefenwahrnehmung.htm>
2. Hamatschek, H.: Konzeption und Entwicklung einer auf 3D-Koordinaten basierenden prototypischen Multimedia Anwendung zur Synchronisation zwischen Raumklangsystemen und computergrafischer Produktion im speziellen Kontext des Themas InnoTP. Schmalkalden: Diploma Thesis, Fachhochschule Schmalkalden (2006)
3. Umble, E.A.: Making it real: the future of stereoscopic 3D film technology. *ACM SIGGRAPH Computer Graphics Quarterly* 40, 1 (2006)
4. InnoTP: Project homepage (Last access: 18.9.2008) (in German), <http://www.innotp.de>
5. IOSONO: Product homepage (Last access: 18.9.2008), <http://www.iosono-sound.de>
6. Sporer, T.: Wave Field Synthesis – Generation and Reproduction of Natural Sound Environments. In: Proc. of the 7th Int. Conference on Digital Audio Effects (DAFx 2004), Naples, Italy (2004)

7. Murray, J.: Did It Make You Cry? Creating Dramatic Agency in Immersive Environments. In: Subsol, G. (ed.) ICVS-VirtStory 2005. LNCS, vol. 3805, pp. 83–94. Springer, Heidelberg (2005)
8. Howard, D., Mabley, E.: *The Tools of Screenwriting. A Writer's Guide to the Craft and Elements of a Screenplay*. St. Martin's Press, New York (1993)
9. Gorsboth, S.: *Prinzipien und Konzepte bei der stereoskopischen Abbildung in virtuellen computergraphischen Umgebungen und deren praktische Umsetzung mittels eines Demonstrators*. Diploma Thesis, Fachhochschule Schmalkalden (2005)

The Functions of Music in Interactive Media

Axel Berndt¹ and Knut Hartmann²

¹ Department of Simulation and Graphics
Otto-von-Guericke University of Magdeburg, Germany
aberndt@isg.cs.uni-magdeburg.de

² University of Applied Sciences Flensburg, Germany
knut.hartmann@fh-flensburg.de

Abstract. In this media theoretical and critical elaboration we intend to close the gap between music in static linear and non-linear interactive media. We will give a brief overview on media music perception and its historical development that enables us to recognize parallels between music in films and games and, thereby, uncover future perspectives of interactive media music. According to this, we will elaborate its narrative functions and potentials in order to widen the scope of the field where music meets interactivity.

Keywords: Interactive Media Music, Perception, History, Narration.

1 Introduction

Amongst all the arts, being interweaved by their cooperative coexistence in multimedia, music is one of the oldest and most frequently underestimated. In combination with image and sound it leads a nebulous shadowy existence, seemingly aimless or at least difficult to explain by terms like *background music*. Its tremendous importance only becomes apparent by muting it and experiencing the disillusioning lack of this unconscious something [1].

It is intended to be perceived subconsciously to directly access the limbic system, by-passing the rational thinking in the cerebrum, and affecting its immanent emotional load [2]. The suggestive power, which is thereby emerged through music, can sensitize the audience to emotional and analogous contents [3]. Furthermore, its combination with other media spans an association space, giving it an amazing clearness of content [4] which the musicologist and linguist Norbert J. Schneider calls a *semantization process* [3]. Music becomes meaningful as a narrative medium that does not even express emotions and mood, but also becomes a means for expression of associations and comments. These possibilities were not available from the very beginning. They had to be discovered and developed over centuries and to be established in the general listening habit.

2 History of Media Music

The linkage of music with other media has a tradition as old as music itself. With text and speech song emerges. In the Greek drama it is linked to stage scenery

and the dramatic plot. Within dance and ballet, music becomes a regulator for movement and gesture. The step from the scenic stage works of the renaissance towards opera interweaves music, plot, and text closer than ever before.

In these forms music is the dominating layer. It is not just restricted to the illustration of textual content, but construes and interprets it. Compared with this, music in drama and stage play has a less organic and intrinsic role but still has its artistic tasks. It prepares the audience for the emotional mood of the scene or summarizes it. It accompanies monolog and dialog and performs punctuation tasks [4].

Film is related to theater. However, the first occurrences of music in films—more precisely silent movies—were evidence of a retrogression of established theatrical traditions and conventions that had to be rediscovered. The primary task of musicians was to drown out the noise of the projector. They were free to play any music that fits reasonably to the mood of the scene. The selection depended on personal favors, mood, musical repertoire and skills. Classical works were performed next to banal pop songs and modern dances. They appeared in smaller and larger fragments and broke off in-between phrases quite often [4].

Obviously, these musics were not composed and thus hardly suitable to accompany the film medium. Adorno and Eisler already pointed out that traditional music structures do not work within the film medium [5]. Music had to become formally more open and unsealed. The composers had to learn to follow rapid short-term scene structures without the time for an extensive exposition or thematic treatment.

Finally, the talkie laid the technical base for the complete synchronization and introduced the layers of speech and sound. The music, so far amongst other things having the task to remedy the spooky aloofness of the silent pictures [5], was now free for a more selective and dramaturgically sophisticated use. The more it was displaced out of the focus of conscious perception towards subconsciousness, the more grew its value and importance for the film as a whole.

Today, music is an inherent part of films and developed a multitude of different forms and aesthetics according to the manifold film genres and functions that it accomplishes. For music in interactive media, especially in computer and video games, a similar development process emerges, giving us the opportunity to discover prospective promising perspectives of interactive media music.

3 History and Future of Interactive Media Music

At first, the unrealistic surreal sound effects of early video games, like *Pong* (by Nolan Bushnell, 1972) or *Super Mario Bros.* (Nintendo, 1985, music and sound effects by Koji Kondo), were a concession to the limited technical possibilities. But soon they established an own self-contained aesthetics for sound design that is closer related to music and cartoonish micky mousing effects than to cinematic sound design. They substitute not just real, but also non-existent sound effects, e.g., gesture illustrations like the upwards sliding glissando that illustrates a jump movement, and motivic cues that reward picking up an item.

Technical restrictions forced the music to be relatively simplistic. In *Asteroids* (Atari, 1979) it consists of only two alternating tones, permanently disturbed by sound effects on the same one and only audio channel. The music of *Space Invaders* (by Toshihiro Nishikado, 1978) is a repetitive sequence of four stepwise descending tones, illustrating the approaching of hostile UFOs. The nearer they come, the more do game-play and musical tempo increase, causing a hectic pace. Karen Collins gives further exemplary discussions of aesthetical aspects of C64 and Atari games music [6].

Space Invaders is one of the first and rare examples from the first decades of games music, that demonstrates its dramaturgic use. The usual musical accompaniment is just a nice-to-have background feature without any narrative functions, as can be observed still today. Most games allow the player to switch it off in the setup menu. Of course, this optional kind of music cannot be of any importance and is deemed to be redundant or at least insignificant. The game must be and is playable and understandable without it. As well as film music did more than half a century before, interactive media music started from the outset apart of all the achievements of hundreds of years of music history.

A turning point was the improvement of sound hardware with more sophisticated synthesis abilities and the introduction of the MIDI standard in 1982. It constitutes a homogeneous musical interface for home computer systems as well as for professional sound studios, enabling musicians to create musical data without programming abilities via graphical high-level interfaces.

Not just the composition became more professional now, but also its performance within the games. One of the highlights of the MIDI and audio based music engines was and still is the *Interactive Music Streaming Engine (iMuse)* of Michael Land and Peter McConnell [7]. It offers a number of automated arrangement techniques, enabling the composer to write music that can be performed and organically adapted to the gameplay without any unmusical breaks or cuts. It was the reversion to the insight that film music made approximately fifty years before: There should be neither an antagonism of character and expression between music and film, nor an indifferent relation [4].

But with the upcoming of the CD-ROM medium in the 1990s interactive media music made a step back. The big memory capacities, now available, led into temptation to use memory intensive audio formats like Wave, MP3 or CD-Audio. These provide high-quality sound but are completely static. Consequently, the indifferent relation between music and interactive scene, the abrupt unmusical cuts within the musical accompaniment are still predominant today.

As well as film music once had to outgrow the structural limitations of traditional music, interactive media music has to find new approaches to coalesce with the interactive medium. Composers became aware of the necessity to compose in a latent way so that asynchronous cuts do not appear too flashy. Or the hard cut is smoothed by a short cross-fade and becomes a still asynchronous soft cut. Others write their music as sequential fragments that can be rearranged in real-time. Newer solutions from the last few years incorporate an extensive use of multi-track arrangement and real-time mixing, an overview is given in [8].

This will pave the way for a more sophisticated use in terms of narration and dramaturgy.

4 Narrative Functions of Music in Interactive Media

In every artistic work its piece-parts and formal aspects are not just present for aesthetic reasons, but to support mediating its content. This is known as the *dialectic unity of form and content*. Hence, music is not included in multi-medial environments just as an end in itself, but performs vital narrative functions.

4.1 The Cinematic Heritage

A very basic and widely accepted high-level classification of the film music's narrative functions is that of Eisenstein et al. [9]. They distinguish between audio-visual *Parallelism*, comprising music that follows and expresses the visual content, and *Counterpoint*, describing music that controverts the scene. This scheme was extended by several musicologists like Pauli [10] and Thiel [11]. They introduced a third category—which Thiel calls *Affirmative Picture Interpretation and Illustration*—comprising music that adds new non-visible content but does not contradict the scene.

This very coarse scheme is hardly suited to distinguish the manifold intentions behind the use of music. That is why Zofia Lissa's very meticulous detailed categorization [4] became important. She distinguishes the following eighteen categories of functions the music could perform (or are related to music):

1. Musical illustration of movement and sounds (known as *Micky Mousing*),
2. Emphasis of movement,
3. Stylizing of real sounds,
4. Representation of locations (geographic, ethnic, social),
5. Representation of time (for historical associations),
6. Deformation of sound (for alienation effects),
7. Comment (audio-visual counterpoint),
8. Source music (diegetic music),
9. Expression of (actor's) emotions,
10. Means of immersion,
11. Symbol (e.g., national anthems),
12. Anticipation of subsequent actions,
13. Enhancement and demarcation of the film's formal structure,
14. Multi-functionality of music (the functions are not mutually excluding),
15. Sound effects (and the mixing with music),
16. Speech/Dialog (e.g., punctuation tasks of music),
17. The function of silence ('The rest belongs to the music as well.' Stefan Zweig),
18. Non-functional aspects (for inner-musical and aesthetic purpose).

Even today, half a century later, Lissa's functions are still up to date and only need to be supplemented by a few new and advanced functions of contemporary cinematic practice. An up-to-date approach is presented by Wingstedt [12].

Emotive Class: emotionalize content and acting;

Informative Class: communication of meaning; communication of values; establishing recognition;

Depictive Class: describing settings; describing physical activity;

Guiding Class: attention guidance; mask (out) unwanted or weak elements;

Temporal Class: provide continuity; define structure and form;

Rhetorical Class: comment, make a statement, judge.

In connection with his temporal class, Wingstedt states furthermore: “In interactive non-linear media, such as computer games, music’s ability to provide continuity is an important quality with strong potential.” [12] Likewise all the musical functions can be transferred to interactive media.

However, the only established functions in games are those providing the feeling of immersion into the scenario and a superficial dramatization of action scenes. Occasionally, some functions of Wingstedt’s informative class appear, in terms of a parallelism based recognition establishing use of the leitmotif technique. Hence, interactive media are far behind their possibilities and even behind ancestors. Especially the contrapuntal functions seem completely unknown to game developers. Even issue-related literature does not invest much attention into a professional music conception [13,14].

4.2 Take a Stand on Interaction!

Since the player is primarily part of the real world he perceives diegetic as well as non-diegetic information, the comment of a non-diegetic contrapuntal music, too. Here a major distinction to linear media emerges: extra-diegetic film music cannot be heard by the actors on screen, thus, has no influence on the plot. But in interactive media the player acts in the virtual world and non-diegetic information can influence the diegesis over him. Interactive media music necessitates a different conception!

It cannot refer only to the virtual scene, surrounding the player, demoting him to an external outsider, but take a stand on himself and his acting. By music’s associative power it is possible to make it clear for him what his actions cause and mean. It can laud him for doing something good and reprove him for bad morally condemnable actions. This educationally and even therapeutic interesting potential lies idle!

Moreover, music can be used as a regulator for the player’s attention, emotional state, and playing behavior by an adaptive musical soundtrack that dynamically reacts and mediate a personalized playing experience. It can be a powerful instrument to ensure and fine-tune the type of a particular scene (action, creepy, sneaky) even if the player acts in a different unintended way. Furthermore, music can be used to influence the player’s decision process by accenting some associations and masking out others.

Till this day this potential is widely unused! Game developers have to become aware of the ontological difference between linear and interactive media in all its facets and apply music as an active essential participant in the multimedial interplay and dialog with the player.

5 Conclusions

This paper presented a media-theoretical attempt towards closing the gap between linear and interactive media music. The analysis of film music history uncovered many parallels to the development of interactive media music (technical limitations in the beginning, rising professionalism, dealing with the problem of synchronization), and indicates promising future perspectives:

The synchronization and organic correlation of musical and interactive elements is one of the most challenging technical problems. New musical structures have to be found to enable flexible adaptive performances. On this technical base interactive media music will rediscover and widen its rich pool of narrative functions for which we have shown that interaction must be an integral part to open up the great potential currently unused.

Acknowledgements. We like to thank Jean Martin and Niels Knolle.

References

1. Jørgensen, K.: On the Functional Aspects of Computer Game Audio. In: *Audio Mostly 2006: A Conf. on Sound in Games*, Piteå, Sweden, Interactive Institute, Sonic Studio Piteå, pp. 48–52 (October 2006)
2. Bullerjahn, C.: *Grundlagen der Wirkung von Filmmusik*. Forum Musikpädagogik, vol. 43. Wissner-Verlag, Augsburg (2001)
3. Schneider, N.J.: *Handbuch Filmmusik I: Musikdramaturgie im neuen Deutschen Film*, 2nd edn. Verlag Ölschläger, Munich (1990)
4. Lissa, Z.: *Ästhetik der Filmmusik*. Henschel, Leipzig (1965)
5. Adorno, T., Eisler, H.: *Composing for the Films*. Oxford University Press, New York (1947)
6. Collins, K.: In the Loop: Creativity and Constraint in 8-bit Video Game Audio. *Twentieth-century music* 4(2), 209–227 (2007)
7. Land, M.Z., McConnell, P.N.: Method and apparatus for dynamically composing music and sound effects using a computer entertainment system. United States Patent 5, 315, 057 (May 1994) (November 1991)
8. Berndt, A., Theisel, H.: Adaptive Musical Expression from Automatic Realtime Orchestration and Performance. In: Spierling, U., Szilas, N. (eds.) *ICIDS 2008*. LNCS, vol. 5334, pp. 132–143. Springer, Heidelberg (2008)
9. Eisenstein, S.M., Pudowkin, W.I., Alexandrow, G.W.: *Manifest zum Tonfilm*. In: Albersmeier, F.Z. (ed.) *Texte zur Theorie des Films*, 3rd edn. Reclam (1998) (1st release 1928)
10. Pauli, H.: *Filmmusik: Ein historisch-kritischer Abriß*. In: Schmidt, H.C. (ed.) *Musik in Massenmedien*, Schott, Mainz, Germany, modified in 1977 and 1981 (1976)
11. Thiel, W.: *Filmmusik in Geschichte und Gegenwart*. Henschelverlag Kunst und Gesellschaft, Berlin (1981)
12. Wingstedt, J.: Narrative functions of film music in a relational perspective. In: *ISME Proceedings—Sound Worlds to Discover*, Spain, International Society for Music Education (2004)
13. Marks, A.: *The Complete Guide to Game Audio*. Osborne/McGraw-Hill, USA (2001)
14. Sanger, G.A.: *The Fat Man on Game Audio: Tasty Morsels of Sonic Goodness*. New Riders (July 2003)

Adaptive Musical Expression from Automatic Realtime Orchestration and Performance

Axel Berndt and Holger Theisel

Department of Simulation and Graphics
Otto-von-Guericke University of Magdeburg
Universitätsplatz 2, 39106 Magdeburg, Germany
{`aberndt,theisel`}@isg.cs.uni-magdeburg.de

Abstract. When scoring an interactive scene or narration music has to follow its events and development organically. Approaches toward musical nonlinearity are needed. Therefore, the means of orchestration and expressive performance provide a big potential that has not been tapped even a little yet. In this paper we will show how to translate them into the interactive context. But to change musical expression it is not feasible to simply switch hard between different instrumentations and performative styles. We introduce and discuss a new style-independent method for organic and musically believable transitions of compositional and performative expression characteristics.

Keywords: Interactive Media Music, Music Adaption, Orchestration, Expressive Performance.

1 Introduction

Music that is applied to different media (e.g., text, dramatic plot, moving image) fulfills a multitude of functions. Especially in the film medium the bandwidth of its narrative functions is extensive: it mediates mood, cross references, emotionalizes, comments, affects the attention of the audience, establishes continuity, structure, and form etc. [1]. The use of music in interactive media, for instance, computer games, artistic and multimedia installations, is aimed at the same tasks. An overview and critical discussion of narrative music concepts for interactive media and stories is given in [2].

But when scoring interactive media a main challenge arises from an ontological differentness. In contrast to the interactive scenario, music is statically predefined, its playback is extensively linear and unable to organically follow short-term changes in the interactive context. It appears as permanently cut fragments without coherence and reveals an indifferent relation to the interactive scene. While film music coped with this well-known excoriated teething troubles half a century ago [3,4,5], interactive media music still suffers from it.

Though music does not have to be linear and musical structure offers a lot of starting points to promote adaptability, few has been done in this area yet. After an overview of approaches to music adaption in Section 2 we will give a

short music theoretical discussion in Section 3 to introduce the musical basis for our approach to organically change musical expression characteristics in realtime during playback in Section 4. Section 5 provides a conclusion.

2 Overview Music Adaption

To link music and interactive scene it is always necessary to gear into the currently playing musical material. According to the type of intervention three different strategies can be differentiated: compositional approaches, generative, and a hybrid combination of both.

The **compositional approaches** leave the process of creating art at the real artist, i.e., the human composer. All musical material is precomposed; realtime adaptations are of an arranging nature. The composer enables this by specific architectural inner musical features. To rearrange the sequential order of musical sections (cf. classic musical dice games [6,7,8,9]) their melodic, harmonic and metric connectivity must be postulated. Fading techniques and the combination of simultaneously playing tracks (as applied in [10,11]) premise that they harmonize tonally and metrically with each other.

By contrast, **generative approaches** take over the composition process and create musical data in realtime. Changes can be induced by the manipulation of parameters that affect the behaviour of functional and algorithmic descriptions [12,13,14] or aspects of the internal representation of contrapuntal styles (e.g., neural networks [15] or knowledge bases [16,17,18]) that regulate the generation process. Other music generators are based on templates (e.g., melodic fragments, harmonic progression patterns) that specify a more or less rough direction of the generation process and constitute another way for interaction [19,20,21].

To combine both, the unsurpassable artistic quality of human composed music and the higher flexibility of generative approaches, the mixture of both, the **hybrid approaches**, pursue the strategy to adapt precomposed pieces of music in realtime to transition seamlessly between the running musical material and that of another piece [22,23,24,25] or to change the expressive character of the one running piece [1].

In this last field the *CMERS* system of Steven R. Livingstone constitutes one of the most sophisticated solutions [26,27]. A number of compositional and performative attributes, namely major/minor mode, average pitch, harmonic complexity, tempo, dynamics, and articulation were mapped onto the emotive load of (western tonal) music which can be changed according to these attributes in order to achieve a different character of expression within the same piece of music. Since emotional changes were less significant when the performance was “too mechanical” [28], it was further “humanized” by the application of Anders Fribergs *KTH* rule system [29,30].

The present paper is engaged in this field, the adaption of performative expression characteristics of a piece, and will elaborate and discuss a new approach

¹ The representation format has to allow such in detail manipulations, as *MIDI*, for instance, does. Wave based formats (Wave, MP3, Ogg Vorbis etc.) are less applicable.

based on musical orchestration principles. The next section will introduce the theoretical background.

3 Musical Issues

Up to now, few attention has been given to aspects of the inner-musical coherence when adapting a piece of music. How to transition different performance styles without conflicting with the work's compositional architecture? How to effect organic changes within the compositional material? These are the most relevant questions and prominent problems in music adaption, since they directly affect the believability and success of any changes. Beyond the matter of emotionality it is also essential for the music's symbolic charge, tonality, and in consequence its transparency, comprehensibility, and psychomotor impact (consider the rhythmic stumbling when variations conflict with the metric order).

In [31] we elaborated a structure protective adaption policy for the creation of transitions between different pieces. These partly apply to inner-musical transitions as well, but need further supplementation and specification for this special situation which is less an instance of a contrapuntal compositional task than a performative and orchestrational one. Hence, musical orchestration principles provide numerous important cues for structure-aware inner-musical transition techniques.

The following subsections will provide an introduction to the topic of orchestration and implicate the role of the expressive performance.

3.1 Orchestration: A Concretion

First of all we have to distinguish the terms *orchestration* and *instrumentation* that may not be confused in our further descriptions. With our definition we follow the one of musicologist Ertuğrul Sevsay [32].

Instrumentation is concerned with the combination of alike or different instruments in order to achieve a desired tonality and timbre. Issues are dynamic balance, timbral contrasts or similarity, articulation, the use of different pitch registers and playing techniques of instruments.

Orchestration deals with aesthetic aspects of a composition's instrumentation in order to create, increase, or decrease contrasts between musical sections to realize or emphasize its musical charge, mood, or general character. Since timbre is easier perceived than formal aspects orchestration is an essential tool for the clarification of musical form and the establishment of transparency. When developing transition techniques for different orchestrations of a composition this has to be kept in mind.

Thus, instrumentation comprises the technical implementation of aesthetic intentions in orchestration. This paper deals with both. Further detailed introductions to instrumentation and orchestration can be found in [32,33,34].

As articulation and dynamics were already mentioned as a matter of instrumentation the performative aspects of music cannot be ignored. The next subsection will focus on the means of an expressive performance.

3.2 Expressive Performance

The expressive performance of a piece of music goes beyond the ‘right notes at the right time’. Each musician adds further performative content to the musical raw material to emphasize artistic intentions and to mediate musical content and form. These performative means of expression are:

Dynamics comprises ‘loudness’ instructions (step dynamics) and continuous changes between different loudness levels (crescendo, decrescendo). Dynamic bows are important means to articulate melodic phrases. Furthermore, an orchestral crescendo is a popular means in film music to lead into dramatic scenes.

Tempo describes the beat count per time unit (e.g., 100 beats per minute). Just like dynamics it features two different shapes, stepwise and continuous (ritardando, accelerando) tempo changes. The active work with subtle acceleration, slowing, and delay resolves the impression of a mechanical performance and is important for the figuration of climaxes, melodic destinations, and structural borders. The work with tempo can influence the perception of the pace of a scene (flowing, hastening, or static).

Emphasis of beats in a bar can be regarded as a micro-dynamical scheme that defines (several) weightings within each bar. A relatively flat scheme can create a majestic or equably fluent pace depending on the underlying tempo and metric density. Strong emphasis differences establish a more rhythmic and agile shape.

Articulation describes the way a note is formed. It can affect its dynamic emphasis (e.g., accent, marcato, sforzato) and its length (e.g., staccato, tenuto, legato). Articulation has a major impact on the character of music; a cantabile melodic line (legato) becomes more snappy with a staccato articulation.

Tuning has to be considered especially for the interpretation of several ethnic musics in proper style. The equally tempered scale is not the only temperament; for medieval music the just temperament is more appropriate; several 20th/21st century composers intentionally put instruments out of tune to achieve a special timbre (e.g., in several cues of the Sauron/Mordor theme in Howard Shore’s *The Lord of the Rings* scores).

Effects can be classified into two categories, technical and articulation effects. The technical effects (e.g., hall, echo, flanging) are created in the sound studio or are a result of the acoustic situation at the location of performance. Acoustics is an important means to mediate breadth, wideness, or intimate nearness. Articulation effects (glissando, triller, bend, fall, vibrato) are performed by the musician. Both fill the performance with life and presence.

Beside these intentionally applied means of expressive performance there is a further unconscious aspect influencing the audible result. Human musicians cannot play and reproduce a piece of music as perfect as a computer. There are little variances in all aspects of the performance, the synchrony of parts, and timbre. These depend highly on the musical context (Is it difficult to play? Was the previous tempo faster or slower than the new one? Does the part play at the borders

of the instruments gamut or on a comfortable level?) and are less significant with professional musicians. However, the problem of *humanization* can be regarded as independent from artistic intentions irrespective of exceptions where the composer plays on (in-)abilities of particular musicians to achieve some uncertain effects. Thus, it will play no further role in the following considerations.

The next section will describe and discuss our approach towards a smooth musical nonlinearity by means of orchestration and expressive performance.

4 Automatic Adaption of Musical Expression

One and the same composition can be orchestrated and performed in very different ways. While the instrumentation and performance data for each version can be created statically in advance by the composer, editor, performer, and/or (at least partly by) a performance system [30,35] the transitions between them are triggered by interaction. They can arise at any position and from very different musical contexts. Creating predefined transitions would considerably raise the practical effort and limit the bandwidth of expression within each version; static transitions are insufficient! The possibilities of human musicians end at the point where transitions have to be created in realtime during playback. In this Section we will describe our approach to this problem of which Figure 1 gives a system overview.

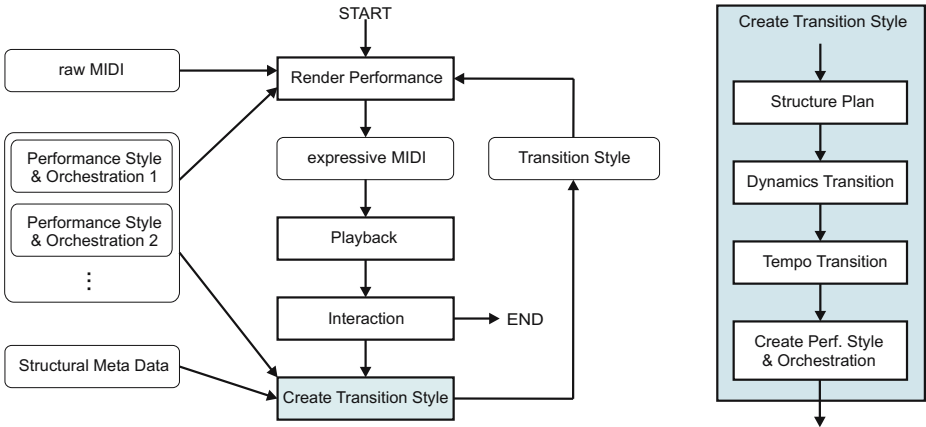


Fig. 1. A system overview of the approach to realtime adaptive orchestration and performance

All implementations were done in C++. As a MIDI-API we used MidiShare [36] which also includes a realtime processing system for MIDI messages. Since the MIDI standard does only support stepwise but no continuous tempo changes, except with lots of very small steps, we implemented an own tempo representation and associated time conversion algorithms for continuous nonlinear tempo progressions.

4.1 Aims

Good transitions melt organically into the music, thus, they seem to be precomposed as well. Therefore, they must not conflict with the musical material. This applies all the more for orchestration and performance, since they are aimed at the reflection and emphasis of this material. Thus, the use of transitional means has to *be justified in the compositional structure*.

Furthermore, a successful transition has to *be efficient* in the way that it varies the running material as little as necessary to get a proper connection to the target material to strengthen the overall coherence.

In the interactive context *latency* is another matter of importance and a fundamental conflict arises: musical change processes need a certain time (the longer the smoother). On the other hand, big latencies are disadvantageous for interactivity, hence, unwanted. Transitions have to take place in the time frame conceded by the application and it must be possible to adapt also a running transition (transition out of a transition) in order to provide the *flexibility* to follow short-term interactions musically.

4.2 Preconditions

An interactive scene is accompanied by different instrumentations of one and the same composition. These emphasize different aspects, e.g., different themes linked to specific elements of the scene while others are masked out, weaved into the accompanying texture, or muted. The same applies to multiple different performance styles that define the way each part should be played. These information are given in separate meta files accompanying the raw MIDI data and are rendered into an expressive MIDI sequence directly before starting the playback.

These different performance styles are stored in an XML data structure comprising information about tempo, dynamics, articulation, and emphasis schemes. Moreover, they declare parts/tracks to be looped, played back only once, or muted in that particular style. The styles do not only describe performative details, but implicate also orchestrational intentions. Different instrumentations, their particular dynamic balance and playing techniques are implemented, e.g., by means of dynamics and articulation.

Per piece of music there is one more file providing information about the time signature and macrostructure locally for each part, or globally for all. Three macrostructural levels are differentiated (listed in decreasing granularity): section structure, phrase structure, and figure structure (motif level). Since, e.g., the last note of a phrase is often already the first note of the next, structural elements are not mutually exclusive, in no strict hierarchical order, and can overlap. This has also been considered in the design of the representation format.

4.3 Approach

To change the style of performance the MIDI data has to be re-rendered. But instead of directly changing the MIDI data a new performance style, the *transition*

style, is created that imports the elapsed beginning of the piece in the current style and the forthcoming rest of the piece in target style. The *transitional section* itself with its certain timely extend is situated intermediate. On the basis of this new style the raw MIDI data can be re-rendered very fast during playback. Since the *transition style* is virtually an autonomous self-contained style it can be the starting point of further transitions, even out of the running one (see the postulate of flexibility in Section 4.1). At the end of the piece, when it is going to be looped, the raw MIDI data is rendered completely in the actual target style.

To create the *transitional section* a systematic approach is implemented. On the basis of musical meta-structure the best suitable start and end points of the transition are determined. According to them the transition of dynamics and tempo is planned in detail. The following explanations will address these aspects. Merely articulation and emphasis scheme are switched directly at the end point, since any interpolations would produce pallid softened results, which is not desired. The switch, in contrast, is a further musically reasonable emphasis on the point where the target style is reached, mediating the impression of deliberate intention.

Structure Planning. Two different application-given latency specifications have to be held: the *begin latency* marks the maximum duration until the transition has to start and the *maximum transition duration* delimitates its timely extent.

Within these latencies the nearest structural borders and their suitability are determined. Structural beginnings are most interesting for new cues of muted parts and for the beginning and end of the transition of already playing parts. Phrase and section beginnings are preferred transition target points (strong ends). If none can be found within the latencies the beginning of the next figure is chosen. It still represents a new content-related beginning, even if not as strong as sections and phrases (weak end). Beginnings are suitable transition target points because from then on a new compositional material starts in the new style; the performative style change reflects compositional structure, thus, appears more appropriate and believable. Structural endings are only interesting to end running parts, since they literally come to an end at these points and should not begin something new.

Such structural borders can not always be found within the latency demands. In this case metrical attributes provide the basis of further planning: If no structure-related starting points can be found within the begin latency the transition begins immediately at the current playback position (weak start). If structural end points are out of range the last emphatic beat within the latency is chosen, which is usually the first beat in the last feasible bar (weak end). Thus, the transition is as long and smooth as possible and its end is placed at an accentuation still providing the feeling of target orientation.

This structural plan forms the basis when applying articulation and emphasis changes and planning the details to transition dynamics and tempo.

Dynamics Transition. Real music instruments usually cannot fade in from or out to zero; normally there is a clearly audible point when they start playing

or stop. New cues for muted parts and endings for running parts by fading, in contrast, produce a more mechanical unnatural impression and should be avoided as far as possible and reserved only as a makeshift in the case of weak borders (this rule does not apply to transitional (de-)crescendos of running parts that play on). If strong entry and exit points were found the parts start/stop directly. Those that stop do additionally lower their volume to $\frac{4}{5}$ of the original level to dynamically trace the gesture of ending; even if this still results in a crescendo it is nonetheless a reserve in the context of the other parts and the musical flow.

For the transition of ongoing parts the dynamics instruction in the target style that is nearest to the structural transition end and inside the latency is chosen as target of the dynamics transition. Its predecessor in the current style marks the starting point. If no dynamics instruction can be found, start and end points are identical to their pendants from the structure plan. The transition between different dynamic levels is generally done with a smooth (de-)crescendo, unless the target instruction defines step dynamics. In this case two different target values are possible, the lower and the upper value of the step. Following the postulate of efficiency (as little change as necessary; see Section 4.1) the value that is nearest to the target value in the current style is chosen.

Tempo Transition. Different tempi are transitioned in the same way as dynamics. But while dynamics can be treated exclusively for each part, tempo must always be global for all parts to keep them synchronous. Therefore, the earliest structural transition start and the latest end point over all parts (unification of the structure plans of each part) is taken as the basis of the tempo transition plan.

4.4 Results

Since a music theoretical analysis provides a substantially more precise armamentarium to discuss the performative and orchestrational quality of transitions than a user study with musically more or less well trained participants we concentrate on a music theoretically shaped discussion in the following.

The approach outlined in this Section opens up a multitude of new possibilities to composers and orchestrators for scoring interactive media. Adaptive expressive performance and orchestration are mighty tools that can shed very different lights on a composition. Combined with the principles of *elastic scoring* (a part can have different instrumentations, expansion and reduction of the number of parts, see [37,38]) and *building set music* or the baroque manner of *rural composition* (optional parts that can be left away, see [10,39]) it is even possible to, beyond timbral changes, flex the compositional form.

All the different versions of a piece of music can be prepared manually by a human artist or (semi-)automatically even during its performance, thus, high artistic quality can be provided. They can be transitioned with much better musical sensibility now. Transitions at strong structural borders (section and phrase borders) are relatively direct, since they fall together with the end and beginning of formally self-contained form elements. At weak positions these changes are done more carefully and unobtrusively.

The independent treatment of tempo and dynamics according to their own structure helps to give the transition a less constructive mechanical impression. It melts more organically with the structural transition plan that only gives a rough orientation. Thereby, e.g., the upbeat character of a figure (its structural borders are not identical with its gestural stress point) can be kept.

However, the current tempo and dynamics transition strategy is to adapt only the region between two instructions, i.e., one (de-)crescendo, ritardando, accelerando or constant area. This approach has its limits: adaptations of very fine granular tempo or dynamics maps produce quite sudden rapid changes that conflict especially with smooth transitions. This can be resolved by carefully interpolating or damping the values of all previous instructions within the transition borders toward the target value of the transition. Therefor, we suggest a potential interpolation to be able to parametrically steer the subtleness of divergence.

These considerations lead to a further question: Do tempo and dynamics transitions need a certain minimum length? In the case of stepwise changes the answer is no. But for continuous changes the overall frequency and amplitude of changes should be taken into consideration. Very different tempi and dynamics need more time to transition smoothly than small nuances.

In the outlined approach we left away the performative aspects of tuning and effects. Thus, following a few remarks on this: It is not suggested to interpolate two different tuning systems during the transition since this produces new and generally unintended tunings that could conflict with compositional intentions. It is recommended to switch them at the structural end of the transition where the target style is achieved, just like emphasis and articulation. Consequently, articulation effects should be handled in the same way, while technical and acoustic effects should be transitioned gradually to add or remove them smoothly to/from the performance.

Aside from dynamics, articulation is another mighty instrument to realize *micro orchestration*, i.e., subtle color variation details, and it can even be “misused” to mute individual notes in order to add or remove, e.g., passing notes or melody variations. Actually, this *mute articulation* provides an interface for nonlinear melodies, articulation effects, and in a harmonic context it can be used to switch between different modes: According to the current performance style a part can play the major or minor third of a triad.

Furthermore, a special performance style that mutes all parts, defines a slow tempo, and lowers the dynamic level can be “misused” in the way that a transition to this style is equal to ending the piece. In case of weak transition end borders it results in a fade-out, but at strong end borders the parts stop directly. Combined with the ritardando and decrescendo a believable finishing gesture is obtained. This works best when all parts have synchronous end points and the transition is not too short to give tempo and dynamics enough time.

In several test scenarios our transition technique has shown a strong stability over the density and completeness of structural meta data. Even in the complete absence of such information the results, although less structure-based, do still

bare a diffuse structure consistency because of the orientation on emphatic stress points and the slight independence of tempo and dynamics planning, which still reflect the compositional structure as described in Section 3.

Since the performative attributes and the way they are handled is largely style-independent, a good stability is also given over different musical styles. Problems may arise with contemporary musical forms. An extensive analysis of stylistic compatibility and the bandwidth of expression that is possible within one composition constitutes a direction of future investigations.

5 Conclusions

The potential of a nonlinear orchestration is largely unexplored, yet. One and the same piece of music can cover a wide range of expression by different orchestration and performance style renderings. We have elaborated and discussed an approach to transition them coherently to resolve the most prominent problem of inconsistency of style changes. Therefore, the structure of the piece of music and its consequential performance attributes played the fundamental role. Hence, the transitions do no longer conflict with the formal structure of the piece and incorporate into it.

By interaction-driven musical change processes a new structural layer emerges that is linked to the interactive context/narration. In the same way as music constitutes the inner structure of films, which is much harder to mediate by visual means, it enriches the interactive medium, helps for a better understanding of narrative coherencies, and mediates nonvisual and nonverbal annotations and connotations. By coherent transition techniques these emanate organically from the musical flow and are more successful because of their stronger believability.

Future work has to musicologically investigate the compositional possibilities opened up by the presented approach. Music/performance production tools for users without XML-programming knowledge (musicians are nonprogrammers in general) are needed. Furthermore, the focus will be on the exploration and development of further approaches to open musical form structures for nonlinearity.

References

1. Wingstedt, J.: Narrative functions of film music in a relational perspective. In: ISME Proceedings—Sound Worlds to Discover, Spain, International Society for Music Education (2004)
2. Berndt, A., Hartmann, K.: The Functions of Music in Interactive Media. In: Spierling, U., Szilas, N. (eds.) ICIDS 2008. LNCS, vol. 5334, pp. 126–131. Springer, Heidelberg (2008)
3. Adorno, T., Eisler, H.: *Composing for the Films*. Oxford University Press, New York (1947)
4. Lissa, Z.: *Ästhetik der Filmmusik*. Henschel, Leipzig (1965)
5. Karlin, F., Wright, R.: *On the Track; A Guide to Contemporary Film Scoring*, 2nd edn. Taylor & Francis, New York (2004)
6. Kirnberger, J.P.: *Der allezeit fertige Polonaisen und Menuetten Komponist (1757)* trans.: *The Ever Ready Composer of Polonaises and Minuets*

7. Mozart, W.A.: Musikalisches Würfelspiel: Anleitung so viel Walzer oder Schleifer mit zwei Würfeln zu componieren ohne musikalisch zu seyn noch von der Composition etwas zu verstehen. Köchel Catalog of Mozart's Work KV1 Appendix 294d or KV6 516f (1787)
8. Joplin, S.: *Melody Dicer*. Carousel Publications, Sparrowbush (1974)
9. Ratai, J.: *Der Jazzwürfel—Ein harmonisches Würfelspiel*. Fraunhofer Institut Medienkommunikation (2005), <http://netzspannung.org/learning/meimus/jazz/>
10. Berndt, A., Hartmann, K., Röber, N., Masuch, M.: Composition and Arrangement Techniques for Music in Interactive Immersive Environments. In: *Audio Mostly 2006: A Conf. on Sound in Games*, Piteå, Sweden, Interactive Institute, Sonic Studio Piteå, pp. 53–59 (October 2006)
11. Wingstedt, J.: REMUPP—a tool for investigating musical narrative functions. In: *Audio Mostly 2006: A Conf. on Sound in Games*, Piteå, Sweden, Interactive Institute, Sonic Studio Piteå, pp. 42–47 (October 2006)
12. Chapel, R.H.: *Realtime Algorithmic Music Systems From Fractals and Chaotic Functions: Towards an Active Musical Instrument*. PhD thesis, University Pompeu Fabra, Department of Technology, Barcelona, Spain (September 2003)
13. Casella, P., Paiva, A.: MAgentA: An Architecture for Real Time Automatic Composition of Background Music. In: de Antonio, A., Aylett, R.S., Ballin, D. (eds.) *IVA 2001. LNCS (LNAI)*, vol. 2190, pp. 224–232. Springer, Heidelberg (2001)
14. Casella, P., Paiva, A.: Mediating Action and Music with Augmented Grammars. In: *4th International Workshop on Intelligent Agents*, pp. 207–211 (2003)
15. Hörnel, D.: *Lernen musikalischer Strukturen und Stile mit neuronalen Netzen*. Shaker Verlag, Aachen (2000)
16. Ebcioğlu, K.: *An Expert System for Harmonizing Chorales in the Style of J. S. Bach*. In: Balaban, M., Ebcioğlu, K., Laske, O. (eds.) *Understanding Music with AI: Perspectives on Music Cognition*, pp. 294–334. AAAI / MIT Press, Cambridge (1992)
17. Schottstaedt, W.: *Automatic Counterpoint*. In: Mathews, M., Pierce, J. (eds.) *Current Directions in Computer Music Research*. MIT Press, Cambridge (1989) (System Development Foundation Benchmark)
18. Löthe, M.: *Ein wissensbasiertes Verfahren zur Komposition von frühklassischen Menuetten*. PhD thesis, University of Stuttgart, Intelligent System Group, Stuttgart, Germany (2003)
19. Pachet, F., Roy, P.: *Musical Harmonization with Constraints: A Survey*. *Constraints Journal* (2001)
20. Pachet, F.: *The Continuator: Musical Interaction with Style*. *Journal of New Music Research* 32(3), 333–341 (2003)
21. Biles, J.A.: *GenJam: Evolutionary Composition Gets a Gig*. In: *Proceedings of the 2002 Conference for Information Technology Curriculum*, Rochester New York, USA (September 2002)
22. Mathews, M.V., Moore, F.R.: GROOVE—A Program to Compose, Store, and Edit Functions of Time. *Communications of the ACM* 13(2), 715–721 (1970)
23. Gartland-Jones, A.: *MusicBlox: A Real-Time Algorithm Composition System Incorporating a Distributed Interactive Genetic Algorithm*. In: Raidl, G.R., Cagnoni, S., Cardalda, J.J.R., Corne, D.W., Gottlieb, J., Guillot, A., Hart, E., Johnson, C.G., Marchiori, E., Meyer, J.-A., Middendorf, M. (eds.) *EvoIASP 2003, EvoWorkshops 2003, EvoSTIM 2003, EvoROB/EvoRobot 2003, EvoCOP 2003, EvoBIO 2003, and EvoMUSART 2003. LNCS*, vol. 2611, pp. 490–501. Springer, Heidelberg (2003)

24. Wooller, R.W., Brown, A.R.: Investigating morphing algorithms for generative music. In: *Third Iteration: Third International Conference on Generative Systems in the Electronic Arts*, Melbourne, Australia (December 2005)
25. Brown, A.R., Wooller, R.W., Kate, T.: The Morphing Table: A collaborative interface for musical interaction. In: Riddel, A., Thorogood, A. (eds.) *Proceedings of the Australasian Computer Music Conference*, pp. 34–39. Australian National University Canberra, Canberra (2007)
26. Livingstone, S.R., Muhlberger, R., Brown, A.R.: Playing with Affect: Music Performance with Awareness of Score and Audience. In: *Australasian Computer Music Conference*, Queensland, Australia (2005)
27. Livingstone, S.R., Brown, A.R.: Dynamic Response: Real-Time Adaptation for Music Emotion. In: *Australasian Conference on Interactive Entertainment*, Sydney, Australia (November 2005)
28. Livingstone, S.R., Muhlberger, R., Brown, A.R., Loch, A.: Controlling Musical Emotionality: An Affective Computational Architecture for Influencing Musical Emotion. *Digital Creativity* 18(1), 43–53 (2007)
29. Friberg, A.: pDM: An Expressive Sequencer with Real-Time Control of the KTH Music-Performance Rules. *Computer Music Journal* 30(1), 37–48 (2006)
30. Friberg, A., Bresin, R., Sundberg, J.: Overview of the KTH Rule System for Musical Performance. *Advances in Cognitive Psychology, Special Issue on Music Performance* 2(2–3), 145–161 (2006)
31. Berndt, A., Hartmann, K.: Strategies for Narrative and Adaptive Game Scoring. In: *Audio Mostly 2007: 2nd Conf. on Interaction with Sound*, Ilmenau, Germany (September 2007)
32. Sevsay, E.: *Handbuch der Instrumentationspraxis*, 1st edn. Bärenreiter, Kassel (2005)
33. Berlioz, H., Strauss, R.: *Grande Traité d'instrumentation et d'orchestration moderne (Treatise on Instrumentation)*. Kalmus Edition, New York, USA (1844/1904) translated english publication (1948)
34. Adler, S.: *The Study of Orchestration*, 3rd edn. Norton & Company, New York (2002)
35. Widmer, G., Goebel, W.: Computational Models of Expressive Music Performance: The State of the Art. *Journal of New Music Research* 33(3), 203–216 (2004)
36. GRAME: *MidiShare Developer Documentation*. Computer Music Research Lab, France. 1.91 edn. (January 2006)
37. Hughes, C.W.: Percy Grainger, Cosmopolitan Composer. *The Musical Quarterly* 23(2), 127–136 (1937)
38. Grainger, P.: *Grainger on Music*. In: Gillies, M., Ross, B.C. (eds.). Oxford University Press, Oxford (1999)
39. Berndt, A.: *Liturgie für Bläser*, 2nd edn. Musikverlag Bruno Uetz, Halberstadt (2008) (transl.: *Liturgy for Brass*)

Narrative Generation for Suspense: Modeling and Evaluation

Yun-Gyung Cheong¹ and R. Michael Young²

¹ Graphics&OS Group, Software Labs, SAIT, CTO, Samsung Electronics CO., LTD.
Giheung-Gu, Gyeonggi-Do, South Korea
yuna.cheong@samsung.com

² Liquid Narrative Group, Computer Science, North Carolina State University
Raleigh, North Carolina, USA
young@csc.ncsu.edu

Abstract. Although suspense contributes significantly to the enjoyment of a narrative by its readers, there has been little research on the automated generation of stories that evoke specific cognitive and affective responses in their readers. The goal of this research is to develop and evaluate a system that produces a narrative designed specifically to evoke suspense from the reader. The system takes as input a plan data structure representing the goals of a storyworld's characters and the actions they perform in pursuit of them. Adapting theories developed by cognitive psychologists, the system uses a plan-based model of narrative comprehension to determine the final content of the story in order to heighten a reader's level of suspense. This paper outlines the various components of the system and describes an empirical evaluation. The evaluation provides strong support for the claim that the system is effective in generating suspenseful stories.

Keywords: Narrative generation, suspense, cognitive model, AI planning.

1 Introduction

Automated story generation has been extensively studied, with applications ranging from computer games to education and training [7; 16; 21]. While a majority of these studies are on automatic generation of logically flawless content, the emotional aspect of storytelling, which is an essential story element for the reader's enjoyment, has received less attention from the interactive storytelling research community. While the computational models of emotion in relation to individual agents have been explored [13; 18], this paper focuses on a rarely explored emotion—suspense—that the audience would feel.

Suspense is the feeling of excitement or anxiety that audience members feel when they are waiting for something to happen and are uncertain about a significant outcome [17; 25; 28]. The significance of suspense in story appreciation has been supported by several studies [1; 3]. In the Brewer and Lichtenstein's study, the participants reported that suspense is cardinal for discerning a story from a mere series of events [3]. Furthermore, the study of viewers' responses to commercials by

Alwitt [1] demonstrates that suspenseful commercials are favored over non-suspenseful commercials. As an effort to explore suspense regarding story structure, Brewer and Lichtenstein [3] claim that affective states in the reader are provoked by arranging the temporal ordering of the events underlying a story world. Their theory explains that suspense could be evoked by presenting the events of a story chronologically to the reader while surprise and curiosity could be caused by hiding a critical fact or event early in the story world and disclosing it later in the text.

This paper presents a computational model of suspense, exploring the concept that a reader’s suspense level is affected by the number of solutions available to the problems faced by a narrative’s protagonists [2; 5; 6; 10; 11; 12; 28]. The reader’s suspense is heightened when undesirable outcomes are likely to happen over preferred outcomes. It is not our intention, however, to deal with the type of suspense that is evoked by visual stimulation such as car chases in film.

Our approach attempts to manipulate the level of suspense experienced by a story’s reader by determining what story elements to tell—that can influence the reader’s narrative comprehension process. To this end, we make use of a computational model of that comprehension process based on evidence from previous psychological studies [2; 10; 12]. To generate suspenseful stories, we set out a basic approach built on a tripartite model, adapted from narrative theory, that involves the following elements: the *fabula*, the *sjuzhet*, and the discourse [22]. A *fabula* is a story world that includes all the events, characters, and situations in a story. In our approach, the *fabula* is represented as a plan structure generated by Crossbow—a hierarchical, partial-order causal link planner based on the Longbow planning system [26]. A *sjuzhet* is a series of events selected from the *fabula* and an ordering over those events indicating the order in which they are to be presented to readers. The final layer, a discourse, can be thought of as the medium of presentation itself (e.g., text, film). Although not directly discussed in this paper, discourse is important for the effective presentation of a story for the reader [4]. Figure 1 presents our three-stage pipelined architecture for story generation as shown, in which Suspenser is situated as a *sjuzhet* generator that determines its contents.




Fig. 1. Tripartite suspense generation model

In this paper, we present Suspenser, a framework that determines narrative contents (i.e., *sjuzhet*) from a given story world (i.e., *fabula*) intended to evoke high level of suspense from the reader as illustrated in Figure 1. We assume in the work described here that the stories we deal with all contain conflict. For example, characters’ individual goals may be negations of each others’, or the plans formed by characters to achieve their goals may interfere with the plans of other characters. While other dramatic devices such as the prolonging of resolution are also useful in crating suspense, we focus here on suspense that arises as a result of users’ consideration of these conflicts and their consequence on the protagonist’s goals.

2 A Computational Model of Narrative Generation for Suspense

Suspenser takes three elements as input: a *fabula*, a given point t in the story plan that corresponds to the point where the reader's suspense is measured, and the story length desired by the system user. The system then determines the *sjuzhet*, the content of the discourse to be used to convey the story up to t to a reader, which enables the reader to infer a minimum number of complete plans for the protagonists' goal, following the psychological research on suspense [10; 12]. In addition, we require that the resulting *sjuzhet* shall be read as a coherent story that represents the input *fabula*.

To produce a *sjuzhet* meeting these requirements, the framework composed of two phases: a skeleton building step and an additional story element identification step. In the skeleton building step, Suspenser identifies the *skeleton* of the *fabula*—a partial plan that specifies its plan steps as a set of core story events that cannot be eliminated without harming the understandability of a story—by rating each individual event's importance based on the event's causal relationship to the protagonists' goals. In the second phase, Suspenser finds actions that can harm the protagonist's goals and tests if the addition of these actions intensifies the reader suspense by modeling the reader's inference process and anticipation of the protagonists' success. The core story events together with harmful actions compose the final content of the *sjuzhet*.

In modeling the reader's inference process and anticipation of the protagonists' success, Suspenser uses Crossbow to model the reader's plan-related reasoning processes. Prior work has provided strong evidence that human task reasoning is closely related to partial-order planning algorithms [19] and that *refinement search* [15], the type of plan construction process performed by Crossbow, can be used as an effective model of the plan reasoning process [27].

2.1 Building the Skeleton

The skeleton building phase determines important events based on the user's knowledge. This step first extracts a series of important events of the story, i.e., a skeleton, and then it tests the skeleton to ensure that its content can be understood as an integral story. To generate a candidate skeleton, the system rates the importance of each event based on a method for extracting important actions that are likely to be included in the story recall, devised by Trabasso et al. [23]. Their approach approximates an individual story event's importance by counting the number of causal relationships with other steps in the narrative and by measuring each event's importance by analyzing its role in a series of actions in a story that are causally related. Adapting their approach, the system computes each step's *importance value* by counting the number of the step's incoming and outgoing causal links of the step and taking into account its role in the plan. For instance, the first action in a story plan and actions that establish the goal state are highly eligible for inclusion in the skeleton. Finally, the top N (the desired story length) events are selected.

Secondly, the system tests whether the skeleton is coherent from the reader's perspective using an algorithm which is a cycle composed of two phases. The first step uses the reasoning algorithm in the reader model to find complete plans to achieve the protagonist's goals that are consistent with the skeleton candidate. If such

a plan is found, the story skeleton is coherent and the program exits. Otherwise, an event in the *fabula* which was not selected as a skeleton with the highest *importance value* is chosen and added to the candidate. Then, the first phase begins again.

Input $\langle t, k, F, W, G, S_K, T_h, R, P, L \rangle$ where

- t is the step where the reader's suspense is measured, k is an integer
- An input *fabula* $F=(SP, B, CL, O)$ where $SP=\{s_1, s_2, \dots, s_i\}$ where s_i is a step, $B=\{b_1, b_2, \dots, b_n\}$ such that b_i is a tuple of $\langle s_i, p_i, v_i \rangle$ where $s_i \in SP$, which means that the plan step s_i binds the parameter p_i to a literal v_i , $CL=\{c_1, c_2, \dots, c_n\}$ such that c_i is a causal link information represented as a tuple $\langle e, s_1, s_2 \rangle$ where e is a condition, and $s_1 \in SP$ and $s_2 \in SP$, which means that plan step s_1 enables the precondition e of s_2 , $O=\{o_1, o_2, \dots, o_n\}$ such that o_i is a temporal constraints represented as a tuple of $\langle s_i, s_j \rangle$ where $s_i \in SP$ and $s_j \in SP$
- $W = \{i_1, i_2, \dots, i_n\}$ such that i_j is a tuple of $\langle s_j, w_j \rangle$ where $s_j \in SP$ when w_j is the importance value of s_j .
- G is the protagonist's goals, T_h is a real number that represents a threshold
- S_K is the portion of the skeleton preceding t
- R is an integer representing the reader's resource bound
- P is a planning algorithm, L is a plan library denoting the reader's knowledge

1 Initialization

- Set $S_T = S_K$, $BSP =$ the portion of SP preceding t
- Set $S = BSP - S_T$, $NZ = \{\}$, $PZ = \{\}$

2 (Construct S_T .) Repeat this step for k times

2.1 If S is empty, return S_T and exit, otherwise

- Pick the action e_S in S generating the highest $h(e_S, F)$. If several candidates are found, non-deterministically select an action with the greatest importance value.
 - If $h(e_S, F) < T_h$, return S_T and exit, otherwise
 - Remove e_S from S .
- Pick the action e_K in S_T with the lowest importance value. If several candidates are found, non-deterministically select an action.
- Replace e_K with e_S
 - If $h(e_K, F) < h(e_S, F)$
 - Construct a partial plan NZ that only contains information in F which has s where $s \in (S_T - e_K + e_S)$.
 - Construct a partial plan PZ that only contains information in F which has s where $s \in S_T$.
 - If $sl(G, NZ, L, P, R) > sl(G, PZ, L, P, R)$
 - $S_T = S_T - e_K + e_S$
 - Add e_K to S
 - Goto step 2.1
 - Add e_S to S .
 - Goto step 2.1

Fig. 2. Content selection for suspenseful stories

2.2 Finding the Additional Story Elements for Suspense

The additional story element identification step constructs the *sjuzhet* (content) evoking the intended suspense level from the reader at t , the target step when the

reader's suspense level is measured. The algorithm in Figure 2 first selects e_S , the action with the greatest *potential suspense*, from the events in the input *fabula* that are not included in the current S_T , where S_T is a series of events to be presented to the reader. If the *potential suspense* of e_S is lower than a predefined threshold T_h , then the program exits and creates a partial plan P composed of the steps in S_T . If the *potential suspense* of e_S passes the threshold, the system chooses the least important action e_K in S_T , and replaces it with the action e_S . Then the system computes the *suspense level* of the newly updated *sjuzhet*. If this substitution lowers the *suspense level* produced from the previous *sjuzhet*, the system brings back the previous value of S_T ; otherwise, the update is maintained. This process repeats until there no candidate is found or for a specified times. When it terminates, the system specifies the content of the output *sjuzhet* as S_T . The next two subsections describe a heuristic function that computes the *suspense level* of a given partial plan, and two heuristics that together compute the *potential suspense* of an action in a plan.

2.2.1 Measuring Suspense Level

In measuring the suspense level on the reader's part, the system follows the notion articulated by Gerrig and Bernardo [12], in which they view an audience as problem-solvers: an audience will feel an increased measure of suspense as the number of options for the protagonist's successful outcome(s) decreases. Adopting these models, we devise Heuristic Function 1 for measuring the *level of suspense*; the function computes the reader's suspense level as the inverse of the number of planned solutions for the protagonists' goal using her reasoning algorithm and her plan library within her reasoning limit. The function sets the minimum level of suspense when no usable solutions are found in her plan space.

Heuristic Function 1 (*Level of suspense*). In the Suspense level function $SL(G, Z, L, P, R)$, G is a set of literals representing the goal of a narrative's protagonist, Z is a partial plan, L is a plan library, P is a planning algorithm, R is an integer representing a reasoning bound, and $success(G, Z, L, P, R)$ returns the number of paths to make G true with given Z and R . $SL(G, Z, L, P, R)$ is set to $(1/success(G, Z, L, P, R))$ when $success(G, Z, L, P, R)$ returns a non-zero value and zero when $success(G, Z, L, P, R)$ returns 0.

2.2.2 Measuring Potential Suspense for an Action

In computing the potential suspense of an action's effect, we consider the action's possible causal relationship to accomplishing the protagonist's goal from the reader's point of view. For example, in a scene in the film *Back to the Future* directed by Robert Zemeckis, the protagonist Marty McFly who just came back to 1985 from 1955 saw Dr. Brown being shot by terrorists. A moment later, however, it was revealed that Dr. Brown was still alive because he was wearing a bullet-proof vest. Although Dr. Brown survived from the shooting after all, the viewers would experience suspense in the shooting scene because they are ignorant of the bullet-proof vest.

In a similar fashion, Heuristic Function 2 computes the *potential suspense* for an action by counting the number of its effects that negate the protagonist's goal and the

number of its effects that unify the goal under the assumption of the audience’s partial knowledge. As an illustration, Figure 3 shows that the action A has an effect $\sim g$, which is the negation of the goal literal g . We call this type of a temporary threat as a *threatening link*, referring to an action’s effect negating another step’s precondition in the plan. In contrast, the suspense creator establishes a *supporting link* when an effect of an action unifies with a precondition of an action in the plan. One effect can have multiple threatening links or supporting links in a single plan. The *potential suspense* of an effect is computed as the supporting link summation subtracted from the threatening link summation as formalized in Heuristic Function 3.




Fig. 3. Threatening links in a story plan. A box represents an action, with its preconditions on the left and effects on the right. Solid arrows denote causal links. Dotted arrows are threatening links which represent an action’s effect negating a precondition of other actions.

Heuristic Function 2 (Potential Suspense of an action). $h(a, p)$ returns the summation of $ps(e, a, p)$ where $ps(e, a, p)$ is the potential suspense of the effect e of the action a in the plan p .

$$h(a, p) = \sum_{e \in effects(a)} ps(e, a, p)$$

Heuristic Function 3 (Potential Suspense of an effect). $ps(e, a, p)$ returns potential suspense of an effect e of an action a in a plan p , which is the summation of the potential threat of the e ’s supporting links subtracted from the summation of the potential threat of the e ’s threatening links as formalized as the following equation.

$$ps(e, a, p) = \sum_{l \in Tlink(e)} \frac{w_t}{dist(d_l, p)} - \sum_{s \in Slink(e)} \frac{w_s}{dist(d_s, p)}$$

Where $Tlink(e)$ returns all the threatening links of an effect e , $Slink(e)$ returns all the supporting links of e , w_t and w_s are coefficients, d_l denotes the destination step of the link l , and $dist(s, p)$ returns a value associated with the causal distance between step a and the goal step of plan p . All scaling factors are constrained to be nonnegative real numbers. In this study, $Dist(a, p)$ returned $d \times (d + 1)$ where d denotes the distance from an action to the goal (i.e., the minimum number of causal links that relate an action to the goal in a plan). The scaling factors w_t and w_s were assigned 7.0 and 2.0 respectively, and the value of the predefined threshold T_h in the algorithm shown in Figure 2 is assigned 0.07 in this study, determined empirically from some informal experiments.

3 Evaluation

While a number of informal experiments and pilot studies have been carried out to evaluate partial implementations of the Suspenser framework [8; 9], this section describes the experiment that we carried out to evaluate the effectiveness of stories that a complete implementation of Suspenser produces in terms of suspense. The hypothesis for our study was to test if there was any association between the story generator type (independent variable) and the suspense level of the stories (dependent variable). To test this hypothesis, the suspense levels among the stories produced by a) Suspenser in high-suspense mode, b) a human author intended to create high suspense, and c) a human author intended to create low suspense were compared to detect a significant difference among them.

3.1 Method

Participants and design. A total of 98 unpaid subjects voluntarily took part in the experiment, ranging in age from 20 to more than 50 years old (42 males, 57 females, and one no response): 72 recruited from NCSU communities including recently graduated under/graduate students across different departments and 26 from internet female technical communities (e.g., Systers.org). All subjects were native-speakers of English. The study utilized a repeated measured between group design: subjects were randomly assigned to one of nine subject groups. These groups were arranged according to a 3×3 Latin Square design to counter-balance the interference from different orderings of stories. From this design, a subject was shown one version of each of the three *fabulas*.

Materials and apparatus. To obtain an input to Suspenser, we ran Crossbow to plan three *fabulas*. The resulting plans consisted of partially ordered 18-20 steps which were manually linearized, and each plan was realized as text using a simple template-matching technique that mapped one plan step into a single sentence. For the study, we prepared a total of nine *sjuzhets*, by generating three *sjuzhets* for each *fabula*—one by Suspenser in high suspense mode and two stories by a human author. One of the *fabulas* and its two *sjuzhets* used in this study are shown in Figure 4. To obtain human generated stories, we recruited one Master's student majoring in English at North Carolina State University, a freelancer writer who had her short story published in a local newspaper. She was presented with texts on sheets that describe the three *fabulas* and their corresponding measurement points. She then was asked to select two series of sentences for each *fabula*: one to arouse high suspense and the other to arouse low suspense from the reader when his suspense level would be measured at a given point in the story. For this study we did not constrain the number of sentences that she selected. As a result, her two versions of a story differed in size within a margin of 2 sentences.

Procedure. Each subject individually participated in the study by accessing a web site. Each subject was presented with three stories and was asked to rate his suspense levels at one point in his reading each of the stories. Each story was presented to the subject sentence by sentence; one page contained only one sentence and a button click led to the

next page. After reading the portion preceding the measurement point displayed on separate pages, the subject was asked to describe his suspense level on a five-point scale basis ranging from “no suspense” to “extremely suspenseful.” After responding to the question, the subject was presented with the second part of the story sentence by sentence, followed by a page asking generic questions about story coherence and enjoyment on a five-scale basis ranging from “not at all” to “strongly agree.”

Background

Sykes is the owner of the Hollywood Theater, which was once prosperous but has now become dilapidated and is in need of major renovations. Sykes has accrued a sizable gambling debt, and with his theater in shambles, he has no means with which to pay it back. He is constantly threatened by his crooked debtors. Janet is a famous actress with dreams of winning an Oscar, an acting award. She is jealous of the actress Agatha, who is her contender for the Oscar this year and also is well-known for her active involvement in charity. Janet knows a number of scoundrels including a guy named Kent, a bomb dealer, and the theater owner Sykes. Agatha is in love with Bill, who serves as a lieutenant in the Los Angeles Police Department's Serious Crime squad. Janet knows that Agatha is planning to go to the Charity Bazaar for the Poor to be held in Hollywood Theater. To ensure that she will win the Oscar, Janet plans to kill Agatha during the charity event.

Storywriter's selection for high suspense effect

Janet and Sykes plan to burn down Sykes' theater to get the insurance money and kill Agatha during the charity bazaar. Janet gives Kent's contact information to Sykes and informs him of Kent's expertise with firebombs. Kent takes a bomb to the Hollywood Theater and meets with Sykes. Sykes purchases the firebomb. Sykes installs the firebomb. Kent informs Bill that Sykes is planning to firebomb his own theater during the charity event. Agatha goes to the theater for the charity event. Sykes sets the timer of the firebomb to explode during the charity event. Sykes switches on the firebomb. Bill searches for the firebomb in the theater. *Bill defuses the firebomb.*

The system's selection for high suspense effect

Kent takes a bomb to the Hollywood Theater and meets with Sykes. Sykes purchases the firebomb. Sykes installs the firebomb. Bill arrests Kent. Kent informs Bill that Sykes is planning to firebomb his own theater during the charity event. Bill releases Kent for his cooperation. Agatha goes to the theater for the charity event. Sykes sets the timer of the firebomb to explode during the charity event. Sykes switches on the firebomb. Bill searches for the firebomb in the theater. *Bill defuses the firebomb. Agatha participates in the charity event.*

Fig. 4. Two *Sjuzhets* produced from *Fabula C*: Italicized sentences are the portion after suspense was measured

3.2 Results

The collected data contained 294 responses from 98 subjects. To detect a significant difference between three story generators, we performed a one-way ANOVA on the collected data using SAS version 9.1.3 SP4. In this analysis, two main effects were examined: the story generator type and the *fabula* type. Each type has three levels.

As shown in Table 1, the data indicated that the story generator type had an effect on the suspense level ($F(2, 285)=4.27$, p value=0.015). The result also shows that the *fabula* type had no effect on suspense. No interaction effect was found between the *fabula* type and the story generator type ($F(4, 285)=0.66$, p value=0.622). Despite the

short sample stories, the subjects rated their experience of suspense is “moderate” (Mean=2.571/5.0, SD=1.059) on a five-point Likert scale. The system performance was superior to the other story generators in the categories of *fabula* B (Mean=2.727, SD=1.126) and *fabula* C (Mean=2.939, SD=1.088).

Table 1. Data for Suspense

Means and standard deviations for suspense in each story generator type (N=98)

Suspenser in the high-suspense mode		Human author for high suspense		Human author for low suspense	
M	SD	M	SD	M	SD
2.704	1.057	2.694	1.049	2.316	1.061

Means and standard deviations for suspense in each story and story generator

Story	Generator	N	M	SD
Fabula A	Suspenser	32	2.438	0.914
	Human-HS	33	2.667	0.890
	Human-LS	33	2.303	1.104
Fabula B	Suspenser	33	2.727	1.126
	Human-HS	32	2.656	1.096
	Human-LS	33	2.394	1.144
Fabula C	Suspenser	33	2.939	1.088
	Human-HS	33	2.758	1.173
	Human-LS	32	2.250	0.950

NOTE: Human-HS denotes the human author’s selection intended to create high suspense and Human-LS denotes the human author’s selection intended to create low suspense.

ANOVA summary table for Suspense

Source	DF	SS	Mean Square	F Value	Pr > F
<i>Fabula</i>	2	1.712	0.857	0.76	0.467
Generator	2	9.571	4.786	4.27	0.015
<i>Fabula</i> *Generator	4	2.954	0.738	0.66	0.622
<i>Error</i>	285	319.760	1.122		

A series of standard one-tailed t-tests were used to compare the performance of the three story generators. The results indicate that the stories produced by the system (Mean=2.704) and the human author intended for high suspense (Mean=2.694) were rated as more suspenseful than the version produced by the human author intended for low suspense (Mean=2.316) with a 99% of confidence (Suspenser vs. Human-LS $t(194)=2.56$, p value=0.006; Human-HS vs. Human-LS $t(194)=2.50$, p value=0.007).

3.3 Discussion

The data clearly show that the story generators had an influence on the amount of suspense that the subjects felt. In particular, the stories produced by Suspenser created stories comparable in suspense to those produced by human authors intended for high suspense effect (Suspenser Mean = 2.704; Human author intended for high suspense Mean = 2.694). The results also show that the difference between the suspense levels felt by the subjects from Suspenser's story for high-suspense and the human author's story for low-suspense was significant with a 99% of confidence.

To test if Suspenser selects appropriate content for the effect of suspense, we investigated the contents of the six *sjuzhets*, and the result indicates that the set of stories for high suspense effect differed in content from the set for low suspense effect. The story created by the system overlapped that created by the human author intended for high suspense in 50%-80% of the total number of story sentences (*fabula* A 50%, *fabula* B 60%, *fabula* C 80%). In contrast, the stories created for high suspense overlapped the story created by the author intended for low suspense in 20%-30% (*fabula* A 20%, *fabula* B 20%, *fabula* C 30%). This means that the story event sets targeting high suspense and the set intended for low-suspense tend to be mutually exclusive, for the story events that the author selected for low suspense were not related to the protagonist's goals. To test if the text quality affected the reader's story comprehension, the subjects' responses to story coherency were also analyzed. The data suggest that the participants evaluated the given stories as relatively coherent (Mean=2.938/5.0, SD=1.031).

4 Conclusion and Future Work

The generation of stories by computers has been the focus of research by computational linguists and AI researchers for several decades. Although a number of approaches have shown promise in their ability to generate narrative, there has been little research on creating stories for an intended emotion. To address this problem, we present a computational model that takes a complete story world and selects contents that can manipulate reader suspense at a specific point in its telling. In determining the contents, this approach gauges the suspense level that a reader would feel by modeling the reader's narrative comprehension using a planning technique. This approach takes as input a partial plan indicating the portion of a story that has been conveyed so far and computes the reader's anticipated suspense level based on the inverse of the number of solution plans that can be found to the protagonist's goals in the space of plans she can consider within her reasoning resources.

To generate a partial plan that maximizes the reader's suspense, the system takes a plan as input and selects a set of core events that have high causal connectivity and that also play an important role in the story. The partial plan then is supplemented by harmful actions (e.g., those that conflict with the protagonist's goals) that intensify the reader's suspense level. The model has been implemented and formally evaluated. The data from the experiments have shown this system to be successful in selecting content that elicits high suspense. In particular, the data show that, in the context of

our experiments, this model was as effective as a human author in generating suspenseful stories.

While the results of this study show that Suspenser was effective in generating suspenseful stories, the design of the experiment does not allow us to point conclusively at single reason for its effectiveness. For instance, the plan representation used in this study did not allow a plan to have conflicting goals; a plan structure used in this research was considered a sound solution plan only when it contains no conflicts. In order to create conflicting situations—critical conditions for suspense—the characters' goals were manually specified to foster a compelling story. As a result, protagonist's and antagonist's plans were often related via causal relationships. A redesign of the experiment to use a more conflict-expressive plan representation is needed to better characterize the contribution of the system in the readers' level of suspense.

We plan to extend this model to interactive environments by expanding previous related work on narrative replanning techniques [14; 20]. Our future work also includes bidirectional interactions among the *fabula*, *sjuzhet*, and discourse layers. For example, the technique of postponing story resolution has often been employed for the effect of suspense in human-authored narratives. This has been computerized in the MINSTREL system by inserting additional events that detail the protagonist's struggles in between the story's climax and its resolution [24]. With a bidirectional interaction model, Suspenser could revise the *fabula* to include auxiliary events that situate the protagonist in a seemingly dangerous position upon the request from the *sjuzhet* generator. Likewise, the *fabula* and the *sjuzhet* could be adjusted upon request from the discourse generator.

Acknowledgments

Partial support for the work described in this paper has been provided by National Science Foundation CAREER award # 0092586 and award 0414722.

References

1. Alwitt, L.F.: Suspense and Advertising Response. *Journal of Consumer Psychology* 12(1), 35–49 (2002)
2. Brewer, W.F.: The Nature of Narrative Suspense and the Problem of Rereading. In: Vorderer, P., Wulff, H.J., Friedrichsen, M. (eds.) *Suspense: Conceptualizations, Theoretical Analyses, and Empirical Explorations*, pp. 107–127. Lawrence Erlbaum Associates, Mahwah (1996)
3. Brewer, W.F., Lichtenstein, E.H.: Stories Are to Entertain: A Structural-Affect Theory of Stories. *Journal of Pragmatics* 6(5-6), 437–486 (1982)
4. Callaway, C.B., Lester, J.C.: Narrative Prose Generation. *Artificial Intelligence* 139, 213–252 (2002)
5. Carroll, N.: Toward a Theory of Film Suspense. *Persistence of Vision* 1, 65–89 (1984)
6. Carroll, N.: The Paradox of Suspense. In: Vorderer, P., Wulff, H.J., Friedrichsen, M. (eds.) *Suspense: Conceptualizations, theoretical analyses and empirical explorations*, pp. 71–92. Lawrence Erlbaum Associates, Mahwah (1996)
7. Cavazza, M., Charles, F., Mead, S.J.: Character-Based Interactive Storytelling. *IEEE Intelligent Systems* 17(4), 17–24 (2002)

8. Cheong, Y., Young, R.M.: A Computational Model of Narrative Generation for Suspense. In: AAAI 2006 Workshop on Computational Aesthetics, pp. 8–15 (2006)
9. Cheong, Y., Young, R.M.: A Framework for Summarizing Game Experiences as Narratives. In: AIIDE 2006, pp. 106–108 (2006)
10. Comisky, P., Bryant, J.: Factors Involved in Generating Suspense. *Human Communication Research* 9(1), 49–58 (1982)
11. de Wied, M.: The Role of Temporal Expectancies in the Production of Film Suspense. *Poetics* 23, 107–123 (1994)
12. Gerrig, R., Bernardo, D.: Readers as Problem-solvers in the Experience of Suspense. *Poetics* 22, 459–472 (1994)
13. Gratch, J., Marsella, S.: Tears and Fears: Modeling Emotions and Emotional Behaviors in Synthetic Agents. In: The 5th International Conference on Autonomous Agents, pp. 278–285 (2001)
14. Harris, J., Young, R.M.: Proactive Mediation in Plan-Based Narrative Environments. In: Panayiotopoulos, T., Gratch, J., Aylett, R.S., Ballin, D., Olivier, P., Rist, T. (eds.) IVA 2005. LNCS, vol. 3661, pp. 292–304. Springer, Heidelberg (2005)
15. Kambhampati, S., Knoblock, C.A., Yang, Q.: Planning as Refinement Search: A Unified Framework for Evaluating Design Tradeoffs in Partial-Order Planning. *Artificial Intelligence* 76(1-2), 167–238 (1995)
16. Mateas, M., Stern, A.: Façade: An Experiment in Building a Fully-Realized Interactive Drama. In: Game Developer's Conference: Game Design Track, San Jose, California (2003)
17. Ortony, A., Clore, G., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge University Press, New York (1988)
18. Pizzi, D., Charles, F., Lugrin, J.-L., Cavazza, M.: Interactive Storytelling with Literary Feelings. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) ACII 2007. LNCS, vol. 4738, pp. 630–641. Springer, Heidelberg (2007)
19. Rattermann, M.J., Spector, L., Grafman, J., Levin, H., Harward, H.: Partial and Total-order Planning: Evidence from Normal and Prefrontally Damaged Populations. *Cognitive Science* 25(6), 941–975 (2002)
20. Riedl, M., Saretto, C.J., Young, R.M.: Managing Interaction between Users and Agents in a Multiagent Storytelling Environment. In: AAMAS 2003, pp. 741–748 (2003)
21. Riedl, M.O., Young, R.M.: An Intent-driven Planner for Multi-agent Story Generation. In: AAMAS 2004, pp. 186–193 (2004)
22. Rimmon-Kenan, S.: *Narrative Fiction: Contemporary Poetics*. Routledge, New York (2002)
23. Trabasso, T., Sperry, L.L.: Causal Relatedness and Importance of Story Events. *Journal of Memory and Language* 24, 595–611 (1985)
24. Turner, S.: *The Creative Process: A Computer Model of Storytelling and Creativity*. Lawrence Erlbaum Associates, Hillsdale (1994)
25. Vorderer, P.: Toward a Psychological Theory of Suspense. In: Vorderer, P., Wulff, H.J., Friedrichsen, M. (eds.) *Suspense: Conceptualizations, Theoretical Analyses, and Empirical Explorations*, pp. 233–254. Lawrence Erlbaum Associates, Mahwah (1996)
26. Young, R.M., Pollack, M.E., Moore, J.D.: Decomposition and causality in partial-order planning. In: AIPS 1994, pp. 188–193 (1994)
27. Young, R.M.: Using Grice's Maxim of Quantity to Select the Content of Plan Descriptions. *Artificial Intelligence* 115, 215–256 (1999)
28. Zillmann, D.: The Psychology of Suspense in Dramatic Exposition. In: Vorderer, P., Wulff, H.J., Friedrichsen, M. (eds.) *Suspense: Conceptualizations, theoretical analyses and empirical explorations*, pp. 199–232. Lawrence Erlbaum Associates, Mahwah (1996)

A Use of Flashback and Foreshadowing for Surprise Arousal in Narrative Using a Plan-Based Approach

Byung-Chull Bae and R. Michael Young

Liquid Narrative Group, Department of Computer Science,
North Carolina State University Raleigh, NC 27695, USA
bbae@ncsu.edu, young@csc.ncsu.edu

Abstract. This paper describes work currently in progress to develop a computational method for generating flashback and foreshadowing, specifically targeted at the evocation of surprise in the reader's mind. Flashback provides a backstory to explain what caused the surprise outcome. Foreshadowing provides an implicit hint about the surprise. Our study focuses on surprise as a cognitive response rather than as an emotional response. The reader's story construction process is simulated by a plan-based reader model that checks unexpectedness and postdictability of the surprise.

Keywords: Story Generation, Narrative, Flashback, Foreshadowing, Prolepsis, Analepsis, Surprise Arousal, Cognitive Interest.

1 Introduction

In a narrative, there exist two different levels of time. One is story time experienced by the characters in the story world. The other is discourse time experienced by readers [5]. The events in the story plane are related temporally and causally based on their natural order of occurrence, but they are often rearranged in the discourse plane intentionally. In other words, authors can let the readers know about some facts in advance or hide some information until a certain point for a dramatic effect. The narrative theorist Genette [8] explains this temporal disparity using the terms *analepsis* and *prolepsis*. *Analepsis* tells (or shows) what has happened in the past with respect to the present. Similarly, *prolepsis* presents what will happen in the future with respect to "now" in the story. The former is like rewinding the story, and the latter is like fastforwarding the story.

In this paper we use the terms *flashback* as an instance of *analepsis* and *foreshadowing* (or *flashforward*) as that of *prolepsis*, following the conventions in cinematic media [2][5][17]. *Foreshadowing* implicitly alludes to a future event in a manner that makes it difficult for the reader to recognize its meaning until the event actually happens. In contrast, *flashforward* explicitly presents a future event in a manner that makes a reader immediately aware of the impending activity¹. The distinction between *foreshadowing* and *flashforward*, however, may not be clear in film narratives

¹ Based on Genette's terminology, *foreshadowing* is an example of *advance mention* and *flashforward* is an example of *advance notice*. The former is implicit; the latter is explicit.

in which flashforward can be partially explicit with help of the camera (e.g., a shot in which a character's face is hidden by manipulation of the camera angle).

Emotions such as suspense, curiosity, and surprise help the readers focus attention on a story, contributing to the readers' sense of satisfaction [1][18]. These emotions, according to structural affect theory [3][4], can be aroused by manipulation of temporal elements in narrative structure. Empirical studies have shown that temporal manipulation of discourse structure can produce different cognitive and emotional responses by influencing the reader's inferences and anticipation [9][12]. The work described here is motivated by this approach to focus on the use of two narrative devices – flashback and foreshadowing – for surprise arousal in narrative generation.

Although flashback and foreshadowing can be important tools for the author to communicate with the readers for a dramatic effect, methods for the automatic generation of these devices have been rarely investigated in previous computer-based story generation systems.

2 Related Work

In 1976, James Meehan developed a storytelling system called *Tale-Spin*, one of the first attempts at automatic story generation. Since *Tale-Spin*, a number of story generators have introduced various ways to build (or to present) a better story in terms of creativeness, story comprehension, interest, etc. This section outlines major previous works closely related to our study: use of flashback and foreshadowing in narrative, surprise as a cognitive response from literature, and the efforts to generate computer-based stories with a dramatic effect.

2.1 Use of Flashback and Foreshadowing in Narrative

Both flashback and foreshadowing are narrative devices that present story events out of temporal order. Flashback describes some past events related to the present; foreshadowing gives allusion (possibly implicit) to some future events. Typically in film media, flashback often functions to provide backstory in support of a main story line² [17], being presented either as a continuous sequence or as a series of scenes showing only the crux of the backstory. By contrast, foreshadowing, as “hints of what is to come” [5], gives only implicit or partial information. If foreshadowing is completely implicit, the reader realizes its meaning only later in retrospect. If it is explicit with partial information, the reader is forced to fill in the information gap in her mental representation of the story. This kind of foreshadowing often serves to focus the reader's attention on a specific event.

In interactive narrative-centered virtual environments, a user's interaction with a virtual character can prompt flashback or foreshadowing (or flashforward). For example, flashback can depict a character's recall of his or her own past events; flashforward can present a character's imagination of possible outcomes [13]. Our approach is different in that we dynamically generate flashback and foreshadowing from the story at the discourse level.

² Flashback can also be used to refer to an entire main story. For example, a narrator can tell a main story as a form of flashback in retrospect, often with the first person prospective. In this study, we focus only on flashback that functions to describe a backstory.

2.2 Postdictable Surprise as a Reader's Emotion as Literary Response

We read books not just to acquire information but also to receive some kind of reward or to stimulate interest through reading [11][18]. According to Oatley [14], a reader's emotions that arise in response to reading can be classified into two types – external and internal. External emotions are evoked as the reader confronts the pattern (i.e., schema or structure) of the narrative; internal emotions are aroused as the reader enters the story world described in the text. Cognitive responses such as curiosity or surprise epitomize external emotions that occur from narrative structure. Empathy with characters in a story is an example of internal emotions. This classification is in accord with Kintsch's distinction between *cognitive interest* and *emotional interest* [11]. The former arises from a well-organized discourse structure; the latter from emotional context in the story. In our present study we focus only on surprise as cognitive interest rather than as an emotional interest or an internal emotion.

Excluding emotional interest, Kintsch introduces the notion of *postdictability* that can contribute to the value of cognitive interest regardless of story types [11]. Postdictability characterizes a story structure in which every part makes sense for the reader as a whole so that she can construct "a coherent macrostructure" with no conflict in retrospect. Surprise or curiosity based on unexpected events, without consideration of postdictability, will produce no interest to the reader because it is likely to fail to contribute to building a coherent story as a whole in the reader's mind. For this reason, surprise associated with unexpected important story outcome should be postdictable [11] [18].

2.3 Computational Model of Story Generation with Dramatic Effects

Our study is primarily motivated by two previous story generation systems: Minstrel [21] and Suspenser [6]. Minstrel employs various creative writing techniques such as generation of suspense, tragedy, and foreshadowing. Particularly, Minstrel aims to create a sense of inevitability and a sense of unity in the story by foreshadowing contrived or unexpected events. As a criterion for contrivance, Minstrel uses a notion of uniqueness (or uncommonness). In other words, a story event is contrived or uncommon if Minstrel has no memory of it. While this heuristic has benefits in terms of domain independence and learning, the selection of uncommon or unexpected events remains a difficult problem to solve.

Like Minstrel, Suspenser creates a story with a dramatic effect, but it focuses only on one cognitive interest: suspense. Suspenser selects important actions for inclusion in a story using a plan-based reader model that measures the suspense level of the reader at a certain point while reading a story. While Suspenser concentrates on the selection of story content to increase suspense level, our system stresses the presentation order of story content.

Based on the efforts made by Minstrel and Suspenser, we focus on building foreshadowing and flashback automatically in computer-based narrative generation using a plan-based approach.

3 System Architecture

In this section we describe Prevoyant, a computational model of flashback and foreshadowing currently under development. Prevoyant will produce as output a story containing structure intended to evoke surprise in the reader's mind. Given a source story described using a plan data structure, Prevoyant determines the content and insertion point for flashback and foreshadowing events. Prevoyant makes use of a reader model which reflects a reader's conception of a story world constructed during reading. The temporally reconstructed story plan is visualized in a 3D virtual world. To visualize the stories Prevoyant produces, we employ Zocalo, web-based intelligent planning services for control of 3D virtual worlds [22]. Fig.1 illustrates the overall input and output process of Prevoyant.

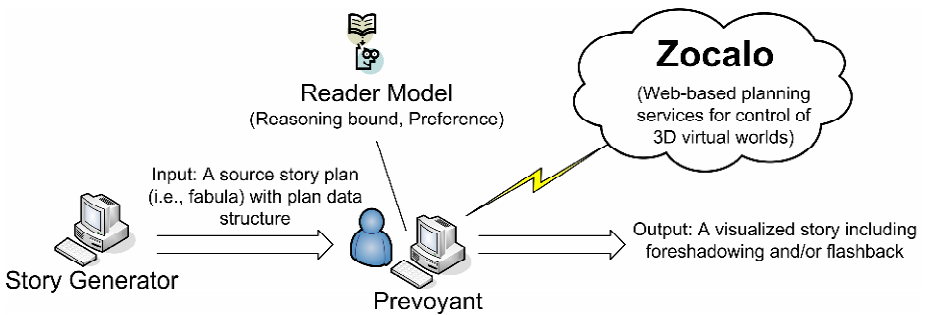


Fig. 1. Input and Output of Prevoyant with the Reader Model and Zocalo Environment

3.1 Generation of Fabula as Source Story

The distinction between *fabula* and *syuzhet*, made by Russian formalism, is a useful concept to understand any narrative element with regard to the ordering of story events [2][8]. The term *fabula* refers to a storyworld in which story events occur in chronological order. The term *syuzhet* characterizes selected contents of the fabula arranged in a certain presentation order, taking into account the reader. To understand a story, the reader constructs her mental representation about a fabula indirectly by drawing inferences about its structure based on an observed *syuzhet* [2]. In our research, a source story (i.e., a fabula in chronological order) is built by Longbow, a discourse planner employing a partial-order causal-link planning algorithm with hierarchical action decomposition [23]. This story is then changed to a *syuzhet* through temporal reconstruction (and possibly through a content selection process).

The plan data structure of the source story created by the Longbow planner includes sets of plan steps, binding constraints, ordering constraints, causal links, and decomposition links. The rich causal relationships in the plan structures created by Longbow correspond closely with the characteristics of the causal networks used by Trabasso and Sperry to represent the relationship between causal relatedness and importance of story events [20]. They claim that the causality among story events is an essential factor for story recall and judgments of important story events [19]. Prevoyant ranks importance of story events based on this causal network story model.

3.2 The Reader Model

Prevoyant's reader model simulates the reasoning process of an *implied reader* that refers to an imaginary, ideal reader as a counterpart to an *implied author* [5][16]. Without the notion of an implied reader, the story may provide readers either with too much explanation or with a sudden change of context [7]. Either can complicate the reader's comprehension process. Prevoyant employs a plan-based reader model using the Longbow planning system, on the basis of a study that has shown that abilities of human planning can be described as those of partial-order planners [15]. Traditionally, AI planners have played a role as a problem solver in task environments. In our approach, we use the Longbow planning system as a reader model in two specific ways: (1) when constructing foreshadowing, (i.e., when a hint about an important story outcome is given), the reader model checks whether or not the reader can infer the outcome as a logical consequence on the basis of her current knowledge; (2) when constructing flashbacks, (i.e., when backstory elements supporting an important story outcome are given), the reader model checks whether or not the reader can comprehend the outcome without any logical conflicts with what she read so far.

The Longbow planning system is based on a domain-independent planning algorithm named DPOCL (Decompositional Partial Order Causal Link) [23], consisting of three components: a *plan-space search* [10] algorithm, a function to represent a user's reasoning resource limits, and a heuristic function used by the planning algorithm to guide the plan-space search process. The process of plan-space search searches through a plan space in which a plan is refined to another plan by fixing a flaw in it. The flaw can be an open precondition (i.e., no actions to achieve this precondition in the plan), a threat (i.e., a step whose effect conflicts with a causal link in the plan), or an abstract step which can be decomposed further. A solution is a complete and consistent plan (i.e., a plan with no open preconditions, no unexpanded abstract steps, no threatened causal links, and no cycles in the temporal ordering constraints). This reader model simulates plot-related inferences performed by the reader. The (possibly partial) plans in the plan space represent completions of the story by the reader. The reasoning bound function defines the limits on the reader's ability to infer plans. The heuristic function characterizes the reader's preference for plans in the planning process such as the preference for content selection. The reader's current knowledge while reading a story is represented by a plan library that defines a set of operators. This reader's plan library characterizes world knowledge and text knowledge based on what she has read so far.

3.3 Prevoyant

Prevoyant aims to create surprise at an important story outcome, which can make the reader more engaged in the story [1]. For the surprise arousal, Prevoyant uses two narrative techniques: foreshadowing and flashback. Foreshadowing provides the reader with (possibly implicit) anticipation³; flashback explains what caused the surprise (i.e., the unexpected outcome) to happen. To meet this end, Prevoyant employs a

³ This kind of foreshadowing is also known as *Chekhov's Gun*, which is a literary technique known from Chekhov's quote: "one must not put a loaded rifle on the stage if no one is thinking of firing it".

generate-and-test design incorporating three major components: the generator, the evaluator, and the implementer. The generator and the evaluator work together to reconstruct a given story based on inferences directed by the reader model. Specifically, the evaluator checks unexpectedness and postdictability in terms of surprise arousal. After the reconstruction of story events is complete, the implementer decides manifestation details based on the specific medium in which the story is being told. The subsequent sections explain these three components in more detail.

3.3.1 Generator and Evaluator

Given in chronological order, events in a source story (i.e., a fabula) are temporally reconstructed by the generator and the evaluator in Prevoyant. The generator creates potential flashbacks and foreshadowings. The evaluator checks unexpectedness for the flashbacks and postdictability as a whole, taking into account the reader model.

3.3.1.1 The Generator: Flashback Selection. Flashback provides the reader with a backstory occurring in the past, typically associated with a character, an object, or an event [17]. The use of flashback, however, should be carefully designed because frequent use of flashback, especially to explain some insignificant elements in the main story, may harm the story momentum.

A narrative structure that evokes surprise is characterized by “sudden presentation of an unexpected outcome”, where expository or initiating events associated with the outcome are presented after the outcome or even omitted [18]. This narrative structure for surprise arousal is similar to that used for curiosity arousal. Both of these structures present unexpected outcomes without their initiating events, but in the two cases the reader’s knowledge about the initiating events is different. If the reader knows that the initiating events are missing or only partially depicted, curiosity occurs; if the reader is not aware of the absence of the initiating events, surprise occurs [1][4][18]. Based on this model, the generator selects flashback events by identifying a Significant Event (SE) and its Initiating Events (IE) in the input story. The algorithm for flashback selection is outlined as follows:

- I. Given an input plan representing a story, identify a set of plan steps that directly connect to goal literals in the goal state, assuming that goal literals are important outcomes in the author-centric story generation system⁴ [1]. Define this set as a candidate SE set.
- II. Consider each SE in the candidate SE set. Let the IE for this set be just a series of events which are causally linked from the initial state to the SE. We say that an IE is *separable* just when the IE can be omitted without causing any open preconditions in the story plan with regard to the events before the SE. Define this pair of the SE and its IE as a separable causal chain. If an SE contains more than one separable causal chain, non-deterministically pick one of these separable causal chains. If there are multiple SEs that have separable causal chains, select the SE with the highest importance rank. The importance is rated in terms of causality, based on the claim that causal connections are more likely to predict recall [20]. The more causal links a candidate SE has, the more significant it is.

⁴ In this paper we assume that all goal literals have equal importance.

III. Define flashback as the IE on the separable causal chain defined in (II), and put its temporal order after the presentation of the SE.

The separability of the IE ensures that the reader does not detect the absence of the IE from the rest of the story until the associated SE happens. Moreover, the generator can select a flashback as a minimum set of the IE with which the reader can comprehend how the surprise (i.e., the unexpected significant event) could happen, minimizing the influence on the story flow. The minimum set of the IE can be determined by appeal to the reader model.

3.3.1.2 The Evaluator: Unexpectedness Check. When the generator identifies a flashback as a set of separable IE, the evaluator confirms the unexpectedness by reasoning about the impact of the omission of the separable IE on the reader's comprehension process. This process is simulated using plan-space search in the reader model. The search process searches for a complete plan to satisfy the preconditions of the SE, without flashback (i.e., the initiating events on a separable causal chain). If a complete plan is found, the evaluator determines expectedness, that is, the surprise fails. In this case, the generate-and-test process iterates and the generator makes another attempt with a different SE in the candidate SE set. Our present study does not measure the level of unexpectedness. More refined algorithm to achieve qualitative unexpectedness will be necessary in the near future. If no complete plan can be found, the evaluator determines unexpectedness.

3.3.1.3 The Generator: Foreshadowing Selection. In the work we describe here, foreshadowing refers to the omitted IE that will be presented after the presentation of the SE as a form of flashback. The use of foreshadowing can enforce postdictability of the omitted events. As mentioned earlier, postdictability is a crucial factor in surprise arousal [11] [18].

Foreshadowing gives hints to the extent that the reader cannot predict its allusion, but can postdict it in retrospect. The implicitness level of foreshadowing will depend on story types or the reader's preferences. In our present study, the criterion for foreshadowing is defined as the least allusion to the omitted IE. As a way of the least allusion, Prevoyant employs two methods of partial flashforward. One is to present a character or an object in an IE in case it is not yet introduced before the SE happens. The other is to present some events in an IE with the removal of important parameters by marking them as hidden. Either should be checked through the reader model whether or not it retains the unexpectedness. While this foreshadowing can be put at multiple temporal positions before the presentation of the SE, it is conventionally placed in the beginning of the story, minimizing the interruption of the story flow.

3.3.1.4 The Evaluator: Postdictability Check. After the selection of foreshadowing and flashback is complete, the evaluator checks postdictability. In present study we assume that surprise is postdictable if there is no conflict among the story elements when the reader reconstructs the whole story in retrospect. Since the flashback and foreshadowing in our approach are based on a story plan that is created from a sound planner, the temporally reconstructed story in retrospect satisfies its postdictability.

3.3.2 The Implementer

Given a temporally reconstructed story plan with flashback and foreshadowing, the implementer manifests it in a specific medium. Through the interaction with Zocalo [22], the implementer converts the story plan into a script that consists of a series of executable actions in a commercial game engine, such as the Unreal Engine by Epic Games.

As a way to handle flashback and foreshadowing at the discourse level (or at the text/media level), the implementer takes advantage of patterns of camera usage typically seen in film. When flashback is connected to a character, associated visual cues are often followed by flashback: the camera tends to show the character in close up as a signal to the viewer, and then follow with flashback scenes. Using a different color filter for shots, e.g., sepia, can also be an indication of flashback.

Unlike flashback, foreshadowing is given without a visual cue since it is implicit in nature. In film narrative, the camera serves as an implicit narrator by virtue of controlling what is seen in a scene (or a shot). Specifically, *hidden* information in the foreshadowing can be efficiently handled by the camera. For example, consider a foreshadowing with a primitive action *Gunshot* (?shooter, ?victim), where ?shooter and ?victim are binding variables. If the variable ?shooter is marked as hidden, the face or identity of the shooter can be concealed through the manipulation of camera angle. If the both variables are supposed to be hidden, the sound of gunfire and the close-up of the gun can be shown without revealing the shooter or the victim.

3.4 The Zocalo Environment

Zocalo is a service-oriented architecture developed by Liquid Narrative group at North Carolina State University, integrating a variety of planning-related software components such as Fletcher (Web Services), Longbow (a Hierarchical Partial-Order Causal-Link Planner), and the Execution Manager (Intelligent controller of commercial game engine environments) [22]. Based on previous efforts developing the Mimesis system [24], Zocalo provides interface services to translate plan steps in the story plan into executable actions in the 3D virtual environments. By incorporating these service-oriented capabilities of Zocalo, Prevoyant creates stories in 3D graphics. Fig.2 illustrates the Zocalo architecture.

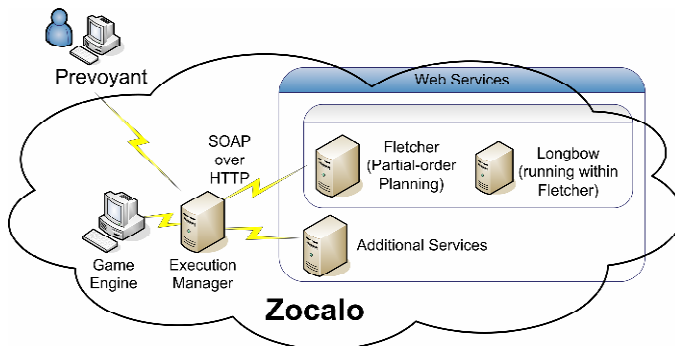


Fig. 2. Architecture of Zocalo, a Web-Based Planning Services

4 Example

In this section, we illustrate how an input *fabula* can be reconstructed to a *syuzhet* with flashback and foreshadowing. As an input *fabula*, we picked an objective story that was previously used for a pilot study measuring suspense in narrative generation [6]. This *fabula* was created by Crossbow, a C# implementation of Longbow, and shares the same plan structure characteristics. The design of the *fabula* – including a plan library, initial state and goal state specifications, etc. – was made by a domain engineer who was not involved in our current effort. This input *fabula* contains a background story represented by ground literals in the initial state and a main story consisting of sentences based on 20 plan steps seen in Fig. 3. The plan steps that comprise a complete story plan were created by the Crossbow planner to satisfy goal literals in the goal state.

For the sake of clarity, we represent causal relationships among plan steps, literals in the initial state, and literals in the goal state as illustrated in Fig.4. Plan steps in the *fabula* are represented by circles. The numbers inside the circles denote the plan step numbers in Fig.3, based on causal and temporal constraints in the *fabula*. In this example, step 1 is the opening step; step 20 is the closing step. The directed arcs between two steps denote causal links from their head-step to tail-step. Specifically, the dotted-line arcs denote the causal links starting from the initial state. The literals in the initial/goal state are represented by rectangles. This *fabula* employs seven ground literals in the initial state, and satisfies seven ground literals in the goal state. The thick solid arrow represents that a goal literal is satisfied by an effect of a plan step (e.g., an effect of the step 4 satisfies the goal literal g1).

Background: The lunatic supervillain known as Jack has been developing biological weapons of devastating proportions. To accomplish the final stages of weapon development, he kidnapped the famous scientist, Dr. Cohen, and brought him to his private fortress on Skeleton Island. Jack expected that the FBI would soon send Smith, their top agent, to rescue Dr. Cohen. To keep the troublesome Smith out of his hair, Jack ordered his own agent, Erica, to monitor Smith and capture him if he is assigned to Dr. Cohen's rescue operation.

Story: (1) Erica installs a wiretap in Smith's home while he is away. (2) Erica eavesdrops on the phone conversation in which Smith is given the order to rescue Dr. Cohen. (3) Erica meets with Smith. (4) Erica tells Smith that her father was kidnapped by Jack and taken to Skeleton Island, and she asks Smith to save her father. (5) Erica gives Smith the blueprints of Jack's fortress, with her father's cell marked. (6) Erica provides Smith with a boat for transportation to Skeleton Island. (7) Before going to the island, Smith hides a diamond in his shoe. (8) Smith goes to the port containing Erica's boat. (9) Smith rides the boat to Skeleton Island. (10) Smith sneaks into the cell marked on the map containing Erica's father. (11) Jack and his guard capture Smith as he enters the cell. (12) The guard disarms Smith. (13) The guard locks Smith into the cell. (14) Smith bribes the guard with the diamond in his shoe. (15) The guard unlocks the door. (16) Smith leaves the cell. (17) Smith sneaks to the lab where Dr. Cohen is captured. (18) Smith fights the guards in the lab. (19) Smith takes Dr. Cohen from the lab. (20) Smith and Dr. Cohen ride the boat to shore.

Fig. 3. Input *Fabula* Story (Adopted from Cheong [6])

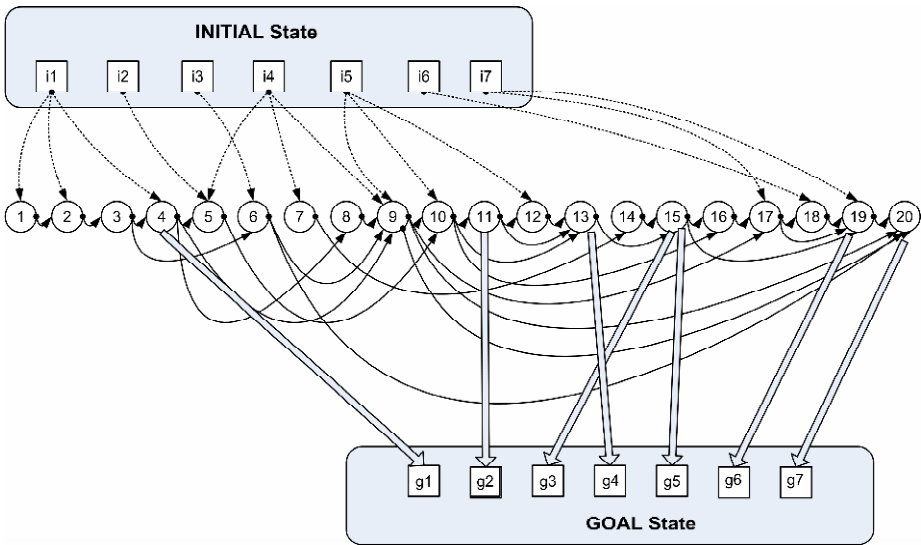


Fig. 4. A Representation of Causal Relationships among Plan Steps, Initial State Literals, and Goal State Literals

According to the algorithm described in the previous section, Prevoyant first selects a candidate Significant Event (SE) set consisting of steps 4, 11, 13, 15, 19, and 20 that directly connect to the goal state. For each candidate SE, Prevoyant checks the possibility of a separable causal chain that will be selected as flashback. For example, step 7 and step 14 can be a separable Initiating Events (IE) as flashback for step 15. The omission of those steps does not cause any open preconditions with regard to the events before step 15. In other words, the reader is not aware of the absence of those events until step 15 is executed. The evaluator confirms its unexpectedness by checking whether or not the reader can infer step 15 without seeing steps 7 and 14. Next, foreshadowing is selected to show partial information of step 7.

In this example, flashback scene consists of two steps: (7) Before going to the island, Smith hides a diamond in his shoe; (14) Smith bribes the guard with the diamond in his shoe. Prevoyant selects step 7, which is farther from step 15 than step 14, for the least allusion. In the story plan, step 7 is instantiated by an action *Hide-Diamond-In-The-Shoe* (?person), where the variable ?person is bound to a constant *Smith*. For implementation of foreshadowing, the camera can show a shot in which someone hides a diamond in his shoe. If step 7 was instantiated with more parameters (e.g., *Hide* (?Smith, ?Diamond, ?Shoes)), depending on the plan library design, the camera could show either a diamond or shoes in the beginning.

5 Discussion and Future Work

This paper describes work currently in progress to develop a computational method for generating flashback and foreshadowing, specifically targeted at the evocation of surprise in the reader's mind. Our study focuses on surprise as a cognitive response,

which is mainly aroused by manipulation of the temporal structure of narrative, rather than as an emotional response. While our present study concentrates on the difference of ordering (of story events) between story time and discourse time, difference of duration or frequency (of story events) between them will be of the same importance [8]. Foreshadowing or flashback, for example, can be presented partially or separately [5]. In that case, the use of foreshadowing and flashback will require careful consideration in order to increase the values of other cognitive interest such as curiosity or suspense at the same time. We plan to take into account the effects of combined cognitive interest (e.g., generation of suspense with surprise outcome or curiosity followed by surprise).

In interactive narrative environments, foreshadowing and flashback will be more closely associated with the cognitive emotion of a story character which a user plays. The content and timing for flashback will be dynamically chosen to reveal the cause of the character's surprise. Moreover, the flashback can show only partial information of the backstory in order to evoke curiosity in the user's mind, which can motivate the user to uncover the whole concealed backstory. Since foreshadowing should be presented prior to the coupled flashback, the system's mediation may occur. In other words, once a foreshadowing is given by predicting a character's surprise, the system will guide the character (that is played by a user) to execute some actions which will lead to the surprise situation. Then the flashback will be given to support the backstory. This kind of flashback can be used to bridge the gap between a main story and its backstories.

References

1. Alwitt, L.: Maintaining Attention to a Narrative Event. In: Shohov, S.P. (ed.) *Advances in Psychology Research*, vol. 18, pp. 99–114 (2002)
2. Bordwell, D.: *Narration in the Fiction Film*. University of Wisconsin Press, Madison (1985)
3. Brewer, W.F., Lichtenstein, E.H.: Event Schemas, Story Schemas, and Story Grammars. In: Long, J., Baddeley, A. (eds.) *Attention and Performance*, vol. 9, pp. 363–279 (1981)
4. Brewer, W.F., Lichtenstein, E.H.: Stories Are to Entertain: A Structural-Affect Theory of Stories. *Journal of Pragmatics* 6, 473–483 (1982)
5. Chatman, S.: *Story and Discourse: Narrative Structure in Fiction and Film*. Cornell University Press, Ithaca (1978)
6. Cheong, Y.: *A Computational Model of Narrative Generation for Suspense*. Ph.D. Dissertation, Department of Computer Science, North Carolina State University (2007)
7. Emmott, C.: *Narrative Comprehension: A Discourse Perspective*. Oxford University Press, New York (1997)
8. Genette, G.: *Narrative Discourse*. Cornell University Press, Ithaca (1980)
9. Hoeken, H., van Vliet, M.: Suspense, Curiosity, and Surprise: How Discourse Structure Influences the Affective and Cognitive Processing of a Story. *Poetics* 26, 277–286 (2000)
10. Kambhampati, S., Knoblock, C.A., Qiang, Y.: Planning as Refinement Search: A Unified Framework for Evaluating Design Tradeoffs in Partial-Order Planning. *Artificial Intelligence* 76, 167–238 (1995)
11. Kintsch, W.: Learning from Text, Levels of Comprehension, or: Why Anyone Would Read a Story Anyway. *Poetics* 9, 87–98 (1980)

12. Levoarto, M.C., Nemesio, A.: Readers' Responses While Reading a Narrative Text. *Empirical Studies of the Arts* 23(1), 19–31 (2005)
13. Marsella, S.C., Johnson, W.L., LaBore, C.: Interactive Pedagogical Drama. In: 4th International Conference on Autonomous Agents, pp. 301–308 (2000)
14. Oatley, K.: A Taxonomy of the Emotions of Literary Response and a Theory of Identification in Fictional Narrative. *Poetics* 23, 53–74 (1994)
15. Ratterman, M.J., Spector, L., Grafman, J., Levin, H., Harward, H.: Partial and Total-Order Planning: Evidence from Normal and Prefrontally Damaged Populations. *Cognitive Science* 25(6), 941–975 (2001)
16. Rimmon-Kenan, S.: *Narrative Fiction*. Routledge, London (2002)
17. Sijll, J.V.: *Cinematic Storytelling: The 100 Most Powerful Film Conventions Every Filmmaker Must Know*. Michael Wiese Productions, Studio City (2005)
18. Tan, E.S.: *Emotion and the Structure of Narrative Film: Film as an Emotion Machine*. L. Erlbaum Associates, Mahwah (1996)
19. Trabasso, T., van den Broek, P.W.: Causal Thinking and the Representation of Narrative Events. *Journal of Memory and Language* 24, 612–630 (1985)
20. Trabasso, T., Sperry, L.: Causal Relatedness and Importance of Story Events. *Journal of Memory and Language* 24, 595–611 (1985)
21. Turner, S.: *The Creative Process: A Computer Model of Storytelling and Creativity*. L. Erlbaum Associates, Hillsdale (1994)
22. Vernieri, T.M.: *A Web Services Approach to Generating and Using Plans in Configurable Execution Environments*, Master Thesis, Department of Computer Science, North Carolina State University (2006)
23. Young, R.M., Pollack, M.E., Moore, J.D.: Decomposition and Causality in Partial Order Planning. In: *Proceedings of the 2nd International Conference on AI and Planning Systems*, pp. 188–193 (1994)
24. Young, R.M., Riedl, M.O., Branly, M., Jhala, A., Martin, R.J., Saretto, C.J.: An Architecture for Integrating Plan-Based Behavior Generation with Interactive Game Environments. *Journal of Game Development* 1(1), 51–70 (2004)

Story Planning with Vignettes: Toward Overcoming the Content Production Bottleneck

Mark O. Riedl and Neha Sugandh

School of Interactive Computing, Georgia Institute of Technology
85 Fifth Street NW, Atlanta, Georgia 30308, USA
{riedl, nsugandh}@cc.gatech.edu

Abstract. Storytelling is prominent part of the daily lives of humans. Entertainers, educators, and trainers often concern themselves with the production of novel stories for entertainment, education, and training. However, it is possible for the consumption of story content by end-users to outpace the rate of production of story content. One solution is to instill greater creativity in computer systems in the form of story generation. We present an incremental advancement to planning-based story generation that increases the space of narratives that can be automatically searched in an attempt to make planning-based story generation more creative. The VB-POCL story planning algorithm implements a form of case-based planning that can incorporate vignettes – plot fragments that are a priori known to be “good” – into a narrative planning process. We show that VB-POCL can generate narratives with favorable structural properties that cannot be generated reliably with previous attempts at planning-based narrative generation.

1 Introduction

Storytelling is a pervasive part of our daily lives and culture. Storytelling is particularly prominent in entertainment, where stories can be viewed as artifacts to be consumed by an audience. Story also plays a role in education and training, where stories and scenarios can be used to illustrate and guide. The production of these artifacts – stories and scenarios – is a primary activity in the entertainment industry and also a significant bottleneck in the educational and training industries. In an “on-demand” society, waiting for periodic updates to serial narratives – weekly television series, movie series, and novels – is not considered ideal. Likewise, players of computer games that rely on stories and quests can complete quests faster than designers can create new quests (for a case study, see [1]). How do we handle the situation in which content consumption outpaces content production? One way to overcome the bottleneck of content production is to instill in a computer system the creative ability to generate new content.

Because of the prevalence of story in non-interactive media such as books and movies, as well as interactive media such as computer games, we concern ourselves with the automated generation of narrative content. The issue is whether an automated story generation system can be considered creative enough or skilled enough to be trusted to produce content – stories – that will be experienced by users. More generally, the output of a creative system, such as an automated story generation system,

must be novel, surprising, and valuable [2]. Whether an artifact is valuable is subjective. For the purposes of this paper, we will consider the minimal requirements for a story artifact to be considered valuable if it (a) meets the intended purpose of its creation and (b) is sufficiently mimetic – appearing to resemble reality, but in a way that it is more aesthetically pleasing than reality. In brief, stories should be novel, but not so novel that they are unrecognizable [3].

In this paper, we describe recent work on planning-based narrative generation. Planning is one model of narrative creation that has been shown to be favorable for story generation (c.f. [4], [5], [6], and [7; 8]). We are considering a model of interactivity in which a user is afforded the ability to specify a set of parameters that abstractly define a desired story. The system then responds to the request by generating a novel story that best meets the given parameters. However, other models of interactivity can also be considered such as *interactive stories*, where a user can participate in a story in real-time. See [9], [10] for techniques for creating interactive experiences by that recursively invoking plan-based narrative generators, and [1] for alternative approach to using planning in interactive narrative system. There are non-planning based approaches to interactive story systems – a non-exhaustive list includes [12], [13], [14], and [15] and derivative works – which are more dependent on hand-authored narrative content and therefore less applicable to the problem of scaling up the pace of content creation. However, further discussion of real-time interactivity is beyond the scope of this paper.

Our goal is to incrementally improve the ability of planning-based story generation by increasing the space of narratives that can be explored algorithmically. We believe that this will provide a capability to address the issue of consumption versus production, and also to create more customized experiences for users under either model of interactivity. While the problem of determining whether a generated story is good is still largely an open problem, we can make claims about the existence of stories with certain properties – properties that could not previously reliably be generated. In particular, we describe the vignette-based partial-order causal link (VB-POCL) narrative planning algorithm. VB-POCL implements a form of case-based planning that can incorporate vignettes – plot fragments that are a priori known to be “good” – into the narrative planning process.

2 Planning Stories

We view story generation as a problem-solving activity where the problem is to create an artifact – a narrative – that achieves particular desired effects on an audience. We favor a general approach where we model the story generation process as planning. One reason for this is that plans are reasonable models of narrative [6]. But also planners “walk the space” of possible narratives in search of a solution that meets certain qualities, making it a good model of creativity in general [2]. This follows from other research efforts modeling story generation as planning (c.f., [5], [7;8]).

The planning process is as follows: a planner chooses an incomplete plan to work on at the fringe of the problem space, and chooses a flaw in that plan to work on, resulting in zero or more new plans in which the flaw is repaired (and often introducing new flaws). These plans become part of the new fringe, and the process is repeated.

One type of flaw pertinent to this work is an *open condition flaw*, which exists when an action in the plan (or the goal state) has a precondition that has not been recognized as being satisfied by the effect of a preceding action (or the initial state). Applying one of the following repair strategies can repair the flaw:

- (i) Selecting an existing action in the plan that has an effect that unifies with the precondition in question.
- (ii) Selecting and instantiating an operator from the domain operator library that has an effect that unifies with the precondition in question.

A planner is non-deterministic, meaning it applies all strategies and then uses a heuristic function to determine which parts of the fringe should be expanded next. There are other types of flaws as well, such as causal threat flaws, which occur when an action threatens to undo the satisfaction of another action's preconditions. Conventional planners assure that plans are sound, meaning that they are guaranteed to execute successfully in the absence of unanticipated changes in the world [16].

However, stories are much more than just ways of achieving an intended outcome in the most efficient manner. Stories should meet the expectations of the audience. This may mean putting in details that are aesthetically pleasing even if they are not strictly necessary. When humans write stories, they call on their lifetime of experiences as a member of culture and society. A computer system that generates stories does not have access to this wealth of information. As a way of mitigating this handicap, a computer system can be provided with a wealth of knowledge in the form of traces of previous problem-solving activities or a library of previous solutions – in this case stories. *Transformational multi-reuse planning* is a form of case-based reasoning in which prior solutions are adapted to new planning problems; systems such as [17] and [18] retrieve and reuse full or portions of old solutions (e.g. plans) to assemble new plans. We adapt the transformational multi-reuse approach and customize it to the particulars of generating stories. However, instead of assuming a knowledge base of complete stories, we bootstrap the planning process with a library of “vignettes” – fragments of stories that capture some particular context.

2.1 Vignettes

We use the term vignette to refer to a fragment of a story that represents a “good” example of a situation and/or context that commonly occurs in stories [19]. For example, a library of vignettes would contain one or more specific instances of bank robberies, betrayals, cons, combat situations, etc. We do not presume to know how these vignettes were created, only that we have the solutions and that they have favorable mimetic qualities. It is important to note that the library contains specific examples of these situations instead of general templates. The implication of the existence of this library is that a story generator does not need to “reinvent the wheel” and thus does not need the specialized knowledge required to be able to create specialized narrative situations. Vignettes are fragments of story structure. How does one know what actions should be included in the vignette and which can be left out? We use the minimal vignette rubric: a *minimal vignette* is one in which removing any one action from the vignette causes it to no longer be considered a good example of the situation and/or context it was meant to represent.

Vignette:	
Steps: 1: Start-Battle (?c1, ?c2, ?place)	Constraints: (character ?c1)
2: Wound (?c1, ?c2)	(character ?c2)
3: Wound (?c1, ?c2)	(stronger ?c2 ?c1)
4: Mortally-Wound (?c2, ?c1)	Variable-constraints: ?c1 ≠ ?c2
5: Die (?c1)	Ordering: 1→2, 2→3, 3→4, 4→5, 4→6
6: End-Battle (?c1, ?c2)	Effects: (battling ?c1 ?c2)
Causation: 1→(battling ?c1 ?c2)→2	(not (battling ?c1 ?c2))
1→(battling ?c1 ?c2)→3	(wounded ?c2)
1→(battling ?c1 ?c2)→4	(mortally-wounded ?c1)
1→(battling ?c1 ?c2)→6	(not (alive ?c1))
4→(mortally-wounded ?c1)→5	

Fig. 1. An example vignette data structure

Computationally, vignettes are stored as plan fragments. As a plan fragment, it is possible that some actions do not have to have all of its preconditions satisfied. This is a way of saying that it is not important how the situation is established or even why, but once the conditions are established certain things should happen. Vignette plan fragments do not reference specific characters, objects, or entities so that a planner can fit the vignette into new story contexts by making appropriate assignments. To ensure illegal or non-sense assignments are not made, co-designation and non-co-designation variable constraints are maintained. Fig. 1 shows an example vignette capturing a very simple combat between two characters where one character (represented by the variable ?c2) is stronger than the other (represented by the variable ?c1). The weaker character wounds the stronger character twice before the stronger character delivers a mortally wounding blow. Finally, the mortally wounded character dies of its wounds. This vignette could be used in any plan in which a character must become wounded, mortally wounded, or dead, or plans in which battles must be started.

2.2 Planning Stories with Vignettes

The Vignette-Based Partial Order Causal Link (VB-POCL) planner is a modification of standard partial order planners to take advantage of the existence of a knowledge-base of vignettes. The VB-POCL planning algorithm is similar to other case-based planners such as [17] and [18] in that it adds a third strategy for repairing open condition flaws:

- (iii) Retrieve and reuse a case that has an action with an effect that unifies with the precondition in question.

Given an action in the plan that has an unsatisfied precondition VB-POCL non-deterministically chooses one of the three above strategies. Strategies (i) and (ii) are performed in the standard way [16]. If strategy (iii) is selected, VB-POCL doesn't immediately repair the flaw. Instead, the plan is annotated with a fit flaw, indicating the plan is to be considered flawed until all actions from the vignette are fitted into the plan. Repairing a fit flaw is a process of selecting an action from the retrieved vignette and adding it to the new plan (or selecting an existing action in the plan that is identical to the selected action to avoid unnecessary action repetition). It may take several invocations of the fitting procedure to completely repair a fit flaw. This may seem more inefficient than just adding all vignette actions to the plan at once.

VB-POCL($P, F, \Lambda_u, \Lambda_v$)

The VB-POCL algorithm takes a plan that is a partial solution to the problem (or the empty plan) P , a set of flaws F evidencing why P cannot be a solution, a library of un-instantiated operators Λ_u that represent templates of actions that characters can take in the world, and a library of vignettes Λ_v .

I. **Termination:** if $F = \emptyset$ and P is sound, return P . Otherwise, fail.

II. Planning:

A. **Goal selection:** Select an open condition flaw $f = \langle s_{need}, p_{need} \rangle$ or a fit flaw $f = \langle P_c, s_c, e_c, p_{need}, s_{need} \rangle$ from F . Let $F' = F \setminus \{f\}$.

B. **Operator selection:** Non-deterministically do one of the following (if valid):

1. **Causal planning (if f is an open condition flaw):** As normal, non-deterministically choosing and instantiating an action from Λ_u or reusing an instantiated operator already in P . If instantiating a new action, add open condition flaws for every precondition of the new action.
2. **Vignette reuse (if f is an open condition flaw):** Non-deterministically retrieve a plan fragment P_c from the vignette library Λ_v such that some step s_c in P_c has an effect e_c that unifies with p . Let f' be a fit flaw such that $f' = \langle P_c, s_c, e_c, p_{need}, s_{need} \rangle$. $F' = F' \cup \{f'\}$. Let P' be a copy of P .
3. **Plan refitting (if f is a fit flaw):** Choose a step s_{next} from P_c that hasn't been chosen before. Let P' be a copy of P . Let s_{add} be s_{next} or a step in P' that is identical to s_{next} (choose non-deterministically). If $s_{add} = s_{next}$, add s_{add} to P' , including relevant causal links, temporal links, and variable bindings. Else, only add relevant causal links, temporal links, and variable bindings from P_c . If $s_{next} = s_c$ then add a causal link from s_{add} to s_{need} in P' . Add necessary open condition flaws to F' for s_{add} for each precondition that will not be satisfied by a causal link in P_c . Let $f' = \langle P_c', s_c, e_c, p_{need}, s_{need} \rangle$ where P_c' is a copy of P_c with s_{next} removed. $F' = F' \cup \{f'\}$.

C. **Threat resolution:** Performed in the standard way (if no resolution exists, backtrack).

III. **Recursive Invocation:** Call VB-POCL($P', F', \Lambda_u, \Lambda_v$).

Fig. 2. The vignette-based planning algorithm

However, there are three advantages to iterative fitting. First, it is easier to recognize and avoid action repetition. Second, it allows for interleaving of repair of other flaws, which can lead to discovery of interesting plans. For example, fitting may lead to the creation of new open condition flaws that in turn are repaired through conventional planning (strategies i and ii) or by retrieving new vignettes (strategy iii). Third, problems in the fitting process can be identified sooner in case the strategy must be abandoned.

The algorithm for VB-POCL is given in Fig. 2. VB-POCL is instantiated with an empty plan P , a set of flaws F , and libraries of un-instantiated action templates and vignettes. Initially, P is an empty plan that only contains information about the initial state – the description of the story world before the story begins – and the outcome state – the description of what the human user wants the story world to be like at the end of the story. VB-POCL selects a flaw. Initially the only flaws are that the outcome state is made up of state propositions that are unsatisfied. As described earlier, open condition flaws, in which an action's precondition (or an outcome state proposition) is unsatisfied, are repaired by three strategies. Strategies (i) and (ii) make up conventional planning (c.f. [16]) and are represented as *causal planning* in Fig. 2. Strategy (iii) is initially handled by the *vignette reuse* portion of the algorithm in Fig. 2. VB-POCL retrieves all vignettes that can satisfy the open condition. How this retrieval happens is relatively simple, but is beyond the scope of this paper. Each successful retrieval creates a branch in the problem space. For each branch, a fit flaw

is created, storing the vignette (P_c), information about the action and precondition that is not satisfied (s_{need} and p), and information about which action in the vignette – called the *satisfier action* – can be used to satisfy the original open condition flow (s_c and e_c). A vignette can be retrieved multiple times with different satisfier actions. The algorithm resolves any causal threats in the normal way (c.f. [16]), and iterates.

When VB-POCL chooses to work on a fit flaw, the *plan refitting* portion of the algorithm is invoked. An action is arbitrarily selected from the vignette and instantiated into the plan (or an identical action that already exists in the plan is chosen). The order doesn't affect the soundness or completeness of the algorithm. All necessary causal links, temporal links, and variable bindings are added to the plan to ensure proper placement and character references of the new or existing action. A special case occurs when the action selected from the vignette is the action that should be used to satisfy the original open condition flow on action s_c . When this happens, an extra causal link is extended from the vignette action to the original action with the unsatisfied precondition. This finally solves the open condition flaw that prompted the vignette retrieval in the first place. To complete the plan refitting process, the action selected from the vignette is removed from the copy of the vignette P_c , and the new plan is annotated with a new fit flaw that has a slightly smaller vignette. The algorithm resolves any causal threats and iterates.

Planning is complete when a plan is found on the fringe that has no flaws.

2.3 Example

To illustrate the VB-POCL planning algorithm, we provide an example of how the planner could use the vignette shown in Fig. 1. Suppose we wanted a story set in J.R.R. Tolkien's Middle Earth. The story world is in the state in which one character, called Enemy, has in his possession a Silmiril – a precious magical stone. The outcome, provided by the human user, is that another character, called Hero, gains possession of the Silmiril. In the remainder of this section, we trace the planning process, describing only one of many possible paths that VB-POCL can follow¹. The planner starts by non-deterministically choosing to satisfy the goal by having Hero take the Silmiril from Enemy. This requires that the Enemy not be alive. The planner could use the vignette from Fig. 1 here by retrieving it and binding Enemy to $?c1$. Note that this strategy will eventually fail because it would require a character stronger than Enemy. Instead the planner chooses to use conventional planning to instantiate an action, *Die-of-Infection*, that causes Enemy to not be alive. This requires that Enemy be superficially wounded. Here VB-POCL retrieves the vignette from Fig. 1 because it has an action that can have the effect (once variables are bound) of causing Enemy to become wounded. Each vignette action is spliced into the new story plan one at a time, using the process of refitting described earlier. Determining where in the plan to splice an action is resolved by repairing causal and temporal inconsistencies (e.g. causal threat flaws). For example, when *Die(Hero)* is spliced into the story plan, it must be temporally ordered after *Take* to avoid inconsistencies; for a character to *Take* an item, the character cannot be dead.

¹ The VB-POCL algorithm follows all possible paths to solving a problem in a best-first manner. We use the term *non-deterministic* to gloss over all the incorrect choices at any given decision-point until it makes a choice that leads to a solution.

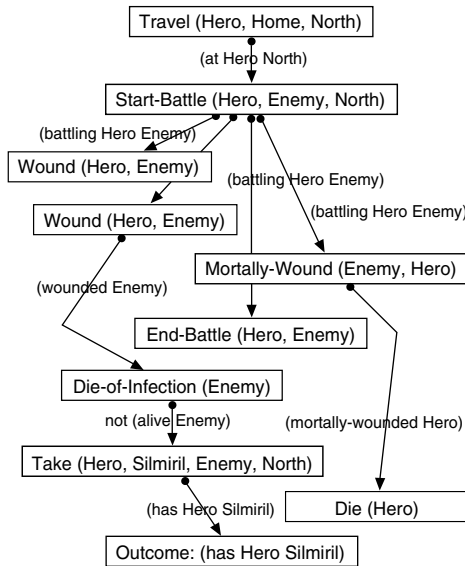


Fig. 3. A story plan generated by VB-POCL

The vignette is fairly self-contained, but the vignette action, *Start-Battle* does require that the planner establish that both Hero and Enemy are at the same place, which in this case the North. This precondition is satisfied in the normal way, by using conventional planning strategies to instantiate an action in which Hero travels to the North (Enemy is already there). The final story plan is shown in Fig. 3. Boxes are actions and arrows represent causal links. A causal link indicates how an effect of one step establishes a world state condition necessary for a precondition of latter steps to be met. For clarity, only some preconditions and causal links on each action are shown.

2.4 Vignette Transformation

VB-POCL assumes a library of vignettes that are already in the domain of the story to be generated. A domain is a set of propositions that describe the world, including characters, and a set of operator templates that described what characters can do and ways in which the world can be changed. In the example, the domain describes characters such as Hero and Enemy and operators such as *Travel* and *Wound*. However, we may want the story planner to have access to vignettes from other domains, especially if our new story is set in an unique and specialized story world domain. We use the term *far transfer* to refer to the process of transferring a vignette from one domain to another.

To engage in far transfer on vignettes, one must first find analogies between domains. Analogy-finding algorithms have been demonstrated to be able to find analogies between stories when stories are pre-existing. Far transfer differs from the problem of finding analogies between stories because there is not a second instance of

a story to compare. Instead, we search for analogies between domains and use that information to translate a known vignette from one domain to another.

The far transfer process is summarized as follows. A *source vignette* is a vignette in an arbitrary domain, called the *source domain*. The *target domain* is the domain of the story to be generated. For each action in the source vignette, the far transfer algorithm searches for an action in the target domain that is most analogical. The search involves a single-elimination tournament where target domain actions compete to be the most analogical according to the Connectionist Analogy Builder (CAB) [20]. The winner of the tournament is the target domain vignette most analogical to the source domain vignette. The result is a mapping of source domain actions to target domain actions that can be used to translate a source vignette into a target domain through substitution.

The far transfer algorithm runs CAB $m*n$ times, where m is the number of actions in the source domain vignette and n is the number of actions in the target domain. The algorithm's complexity is subsumed by the complexity of CAB, which is NP-complete.

Translated vignettes may have gaps where translation is not perfect. This is not a problem because the VB-POCL will recognize this and fill in the gaps via planning. Applying this process to all vignettes in a library results in a new library in which all vignettes are in the proper domain. See [19] for a detailed description of far transfer.

2.5 Discussion

One of the interesting properties of VB-POCL is that vignette retrieval can result in story plans in which there are actions that are not causally relevant to the outcome. Trabasso and van den Broek [21] refers to actions that are causally irrelevant to the outcome as dead-ends. In the example above, the causal chain involving Enemy mortally wounding Hero and then Hero dying appears to be a dead-end because those actions do not contribute to Hero acquiring the Silmiril. Dead-ends are not remembered as well as actions that are causally relevant to the outcome [21], suggesting that dead-ends should be avoided. A battle in which a single wound was inflicted on Enemy would have sufficed, and this is what planners such as [16] and [7] would have settled on.

Human authors regularly include dead-end events in stories suggesting some importance to dead-ends. We hypothesize that there are certain mimetic requirements to be met in any story and that dead-ends can serve this purpose. For example, we assume that a combat scenario in which many blows of varying strengths are exchanged is more interesting than a combat in which a single blow is dealt. Interestingly, what may be a dead-end causal chain to the story planner may not be considered a dead-end by a human reader, and vice versa. That is, the reader may interpret the example story as a tragedy and consider the death of Hero as one of two primary causal chains, whereas the planner's representation contains only one causal chain that leads to the human user's imposed outcome (Hero has the Silmiril). More research needs to be done to create intelligent heuristics to recognize when dead-ends (from the planner's perspective) are favorable, tolerable, or damaging.

VB-POCL is capable of finding stories that other causal-planning based story generation techniques are not able to find. Specifically, these are stories in which some

actions are not strictly necessary for causal achievement of some human-specified outcome state. As noted in [8], expanding the space that can be explored provides an opportunity to find more solutions that are valuable. However, one could claim that some – or all – of the creativity occurred in the process of transforming vignettes, executed prior to generation. As a first step toward improving the ability of planning-based story generation to reliably produce valuable, mimetic stories, the VB-POCL algorithm provides the technical capability of searching a large space of possible solutions. Future work requires strategies for controlling the search space exploration, including heuristics for ranking solution “goodness.” That is, VB-POCL currently has no understanding of how multiple vignettes add or detract from each other or the overall quality of the story being generated.

On a practical note, planning stories with vignettes is a way to increase the average length of stories that can be generated. Ideally, a planner should only have to make $O(n)$ decisions where n is the length of the plan generated. In practice, planners backtrack, meaning that they spend time generating action sequences that do not pan out and must return to an earlier decision point. Any effort spent on a line of reasoning that does not pan out is wasted effort. In the worst case, a planner must consider all ways of making every decision ($O(b^n)$ where b is the number of ways a decision can be made, and n is the length of the solution [16]). Vignettes, when selected, guide the process of adding actions to the story plan, offering up actions in hand-coded sequences that are less likely to result in backtracking than if every action must be chosen independently. Of course, VB-POCL can interleave multiple vignettes during which time new issues that cause backtracking can arise; this is the price of flexibility. Future work is needed to develop powerful heuristic functions that can help VB-POCL discriminate between vignettes when more than one can be applied to an open condition flaw. The practical result of less backtracking is that more time can be spent of fruitful action sequences, potentially allowing for longer plans to be created in less time.

3 Related Work

Search based narrative generation approaches include Tale-Spin [22], which uses a simulation-like approach, modeling the goals of story world characters and applying inference to determine what characters should do. Dehn [4] argues that a story generation system should satisfy the goals of the human user. That is, what outcome does the user want to see? The Universe system [5] uses means-ends planning to generate an episode of a story that achieves a user’s desired outcome for the episode. More recent work on narrative generation attempts to balance between character goals and human user goals [7]. Further work on story planning addresses expanding the space of stories that could be searched [8].

Case-based reasoning (c.f. [23]) has been found to be related to creativity [2; 24]. Several approaches to narrative generation use case-based reasoning. Minstrel [25] implements a model of cognitive creativity based on routines for transforming old stories into new stories in new domains. ProtoPropp [26] uses case-based reasoning to generate novel folk tales from an ontological case base of existing Proppian stories. Mexica [27] uses examples of prior stories to propose plot points and then applies

means-ends planning to fill in missing details. VB-POCL is an attempt to harness the power of search-based generation and case-based creativity in a formalized causal planning framework.

VB-POCL is a variation on case-based reasoning. Case-based reasoners typically engage in four processes: *retrieve*, *reuse*, *revise*, and *retain* [23]. Transformational multi-reuse planners attempt to reuse components of solutions to similar problems to solve new problems, thus possibly invoking retrieve, reuse, and revise processes more than once. VB-POCL is most similar to [17] and [18], but differs from them and all other case-based planners in the following ways. First, vignettes are not complete solutions to previously solved problems; a vignette is not a case. But vignettes are used like a case. Regarding VB-POCL functionality, the VB-POCL retrieval process retrieves all vignettes that can conceivably be used to satisfy an open condition. Typically, the cost of retrieval and reuse is very high so a system must deliberate about the cost tradeoff of standard planning versus retrieval and reuse. Trying all vignettes that meet the requirements for retrieval is not practical if vignettes require extensive modifications for reuse. VB-POCL assumes that vignettes in the library are in the domain of the story being generated and thus do not require extensive effort for reuse. An offline algorithm – summarized in Section 2.4 and described in detail in [19] – is used to transform all vignettes from arbitrary domains into the domain of the new story to be generated. That is, many of the computationally intensive aspects of *reuse* occurs offline and thus fitting a vignette into a plan is trivial. Second, vignettes don't require modification because vignettes are *minimal*. Reuse is performed by blindly inserting all actions in a retrieved vignette into the plan; VB-POCL does not need to make hard decisions about which actions should be kept and which actions should be discarded. Finally, because vignettes are not cases, VB-POCL does not reincorporate (e.g. retain) its solution story plans back into the knowledge base. That is, VB-POCL does not attempt to learn to solve problems from past examples. One of the interesting properties of transformational multi-reuse planning algorithms such as [17] and [18] is that they can operate when there are no applicable cases available; the algorithm can fall back on conventional planning. VB-POCL shares this trait, but unlike transformational multi-reuse planners VB-POCL is also complete, meaning it can find all solutions that exist (the proof is beyond the scope of the paper).

It may be possible to use hierarchical task network (HTN) planners [28] or a decompositional planner such as DPOCL [29] to achieve similar effects as VB-POCL. However, using HTNs or other decompositional techniques to generate story requires reasoning at higher levels of abstraction than the action (or event), and this introduces potentially rigid top-down structuring of plot that can limit opportunistic discovery such as in [7; 8]. Further, vignettes can potentially come from many sources, which may or may not be accompanied by abstract context information.

There are many similarities between VB-POCL and macro-operator planners. Indeed, VB-POCL's vignette retrieval process can be considered analogous to selecting a virtual macro-operator. Macro-operator planners transform an action space into a more compact action space for efficiency gains by learning to group primitive actions that occur together frequently into abstract operators [30]. However, VB-POCL doesn't learn vignettes. Indeed, vignettes often contain action sequences that *cannot* be found by the planner. Further, VB-POCL needs to operate in the primitive action space so that vignettes can be spliced together or so that VB-POCL can use conventional planning techniques to discover new action sequences.

4 Conclusions

VB-POCL is a planning algorithm that extends conventional planning algorithms (e.g. [16]) to make it more applicable to narrative creation. Specifically, VB-POCL extends the conventional planning algorithm by retrieving and reusing vignettes. This is a strategy for tapping into the experiences of other presumably expert story authors. VB-POCL shares many similarities with case-based planners such as [17] and especially [18]. However, by treating plans as narratives – that is, the plan is not a schedule of actions to be executed for goal attainment but a description of events that lead to an outcome – we are able to simplify the case (vignette) reuse problem by assuming that our library of vignettes includes only minimal vignettes.

VB-POCL can explore a greater space of stories because it can consider story plans that have action that are not causally necessary to reach some given outcome. We believe that some of these stories will be more valuable because of the mimetic qualities of the vignettes and the potential for these stories to possess both global novelty and localized familiarity. While we have not yet performed an evaluation of VB-POCL, we believe that this is a step towards instilling computer systems with the ability to assume responsibility for story content creation. This can be an important for application areas where content creation is a bottleneck and it is possible for the pace of content consumption to overtake the pace of content production.

References

1. Morningstar, C., Farmer, F.R.: *The Lessons of Lucasfilm's Habitat*. In: Benedikt, M. (ed.) *Cyberspace: First Steps*. MIT Press, Cambridge (1991)
2. Boden, M.: *The Creative Mind: Myths and Mechanisms*, 2nd edn. Routledge (2004)
3. Giora, R.: *On Our Mind: Salience, Context, and Figurative Language*. Oxford University Press, Oxford (2003)
4. Dehn, N.: *Story Generation after Tale-Spin*. In: 7th International Joint Conference on Artificial Intelligence (1981)
5. Lebowitz, M.: *Story-Telling as Planning and Learning*. *Poetics* 14, 483–502 (1985)
6. Young, R.M.: *Notes on the Use of Plan Structures in the Creation of Interactive Plot*. In: AAAI Fall Symposium on Narrative Intelligence (1999)
7. Riedl, M.O., Young, R.M.: *An Intent-Driven Planner for Multi-Agent Story Generation*. In: 3rd International Joint Conf. on Autonomous Agents and Multi Agent Systems (2004)
8. Riedl, M.O., Young, R.M.: *Story Planning as Exploratory Creativity: Techniques for Expanding the Narrative Search Space*. *New Generation Computing* 24(3), 303–323 (2006)
9. Young, R.M., Riedl, M.O., Branly, M., Jhala, A., Martin, R.J., Saretto, C.J.: *An Architecture for Integrating Plan-Based Behavior Generation with Interactive Game Environments*. *Journal of Game Development* 1, 50–70 (2004)
10. Riedl, M.O., Stern, A., Dini, D., Alderman, J.: *Dynamic Experience Management in virtual Worlds for Entertainment, Education, and Training*. *International Transactions on Systems Science and Applications* 4 (2008)
11. Barber, H.M., Kudenko, D.: *Dynamic Generation of Dilemma-based Interactive Narratives*. In: 3rd AI and Interactive Digital Entertainment Conference (2007)
12. Magerko, B.: *Evaluating Preemptive Story Direction in the Interactive Drama Architecture*. *Journal of Game Development* 2 (2007)

13. Mateas, M., Stern, A.: Architecture, Authorial Idioms and Early Observations of the Interactive Drama Façade, Technical Report, CMU-CS-02-198, School of Computer Science, Carnegie Mellon University (2002)
14. Si, M., Marsella, S., Pynadath, D.V.: Thespian: Using Multi-Agent Fitting to Craft Interactive Drama. In: 4th International Joint Conference on Autonomous Agents and Multi Agent Systems (2005)
15. Weyhrauch, P.: Guiding Interactive Fiction. Ph.D. dissertation, Carnegie Mellon University (1997)
16. Weld, D.: An Introduction to Least Commitment Planning. *AI Magazine* 15(4), 27–61 (1994)
17. Francis, A.G., Ram, A.: A Domain-Independent Algorithm for Multi-Plan Adaptation and Merging in Least-Commitment Planners. In: AAAI Fall Symposium on Adaptation of Knowledge for Reuse (1995)
18. Britanik, J., Marefat, M.: CBPOP: A Domain-Independent Multi-Case Reuse Planner. *Computational Intelligence* 20 (2004)
19. Riedl, M.O., León, C.: Toward Vignette-Based Story Generation for Drama Management Systems. In: Workshop on Integrating Technologies for Interactive Story (2008)
20. Larkey, L.B., Love, B.C.: CAB: Connectionist Analogy Builder. *Cognitive Science* 27, 781–794 (2003)
21. Trabasso, T., van den Broek, P.: Causal Thinking and the Representation of Narrative Events. *Journal of Memory and Language* 24, 612–630 (1985).
22. Meehan, J.: The Metanovel: Writing Stories by Computer. Ph.D. Dissertation, Yale University (1976)
23. Kolodner, J.: Case-Based Reasoning. Morgan Kaufmann Publishers, San Francisco (1993)
24. Kolodner, J.: Understanding Creativity: A Case-Based Approach. In: Wess, S., Richter, M., Althoff, K.-D. (eds.) EWCBR 1993. LNCS, vol. 837, pp. 3–20. Springer, Heidelberg (1994)
25. Turner, S.: MINSTREL: A computer model of creativity and storytelling. Ph.D. dissertation, University of California, Los Angeles (1992)
26. Gervás, P., Díaz-Agudo, B., Peinado, F., Hervás, R.: Story Plot Generation Based on CBR. *Journal of Knowledge-Based Systems* 18(4-5), 235–242 (2006)
27. Pérez y Pérez, R., Sharples, M.: Mexica: A Computer Model of a Cognitive Account of Creative Writing. *Journal of Experimental and Theoretical Artificial Intelligence* 13(2), 119–139 (2001)
28. Sacerdoti, E.D.: A Structure for Plans and Behavior. Elsevier, New York (1977)
29. Young, R.M., Pollack, M., Moore, J.: Decomposition and Causality in Partial-Order Planning. In: 2nd International Conference on AI and Planning Systems (1994)
30. Gratch, J., DeJong, G.: A Framework of Simplifications in Learning to Plan. In: 1st International Conference on Artificial Intelligence Planning Systems (1992)

Schemas in Directed Emergent Drama

Maria Arinbjarnar and Daniel Kudenko

Department of Computer Science
The University of York
Heslington, YO10 5DD, York, UK
maria@cs.york.ac.uk, kudenko@cs.york.ac.uk

Abstract. A common problem in creating interactive drama is the authoring bottleneck. If pre-authored stories are directly incorporated into an interactive virtual environment then there is a need to consider all possible interactions and story twists, which for a sizeable drama is infeasible. One proposed solution to this problem is to use search and planning algorithms along with narrative structures. This leads to a huge state space and planning becomes intractable for real-time solutions. A way to address this problem is to distribute the story planning to autonomous characters so that the drama emerges from their interactions. However without predefined structure or directives the drama is unlikely to emerge into the intended story or even the intended genre. We propose to divide the drama into narrative episodes which we call schemas. Schemas are used by a director and a set of actors to structure the drama so that it emerges into a fully developed drama. The schemas are pre-authored in an abstract way such that they can be deployed multiple times in the same drama, which removes the authoring bottleneck. In this article, we define the structure of the schemas and how the director and actors use schemas in Directed Emergent Drama (DED).

1 Introduction

Creating a truly interactive drama is a difficult challenge. The response of actors needs to be believable, i.e. it needs to be influenced by that which has transpired immediately before and the character's state of mind as well as the emergent drama as a whole. At the same time the emergent drama needs to conform to a specific genre and to a dramatic arc with the expected rise and fall in suspense for the player to be fully immersed in the ongoing drama.

We use a central agent, the director, that directs – but does not prescribe – actions of a set of autonomous actors. By the term "directs" we mean that the director is giving general directions in the form of schemas that guide the actors in developing the drama and engage the player in a Directed Emergent Drama (DED). The director chooses schemas dependent on how the drama plays out and the actions of the player, rather than following a pre-authored script. The schemas are not sequentially ordered, each actor will be in more than one schema at any given time and some schemas are deployed multiple times during each drama. The actor's goal is to entertain and engage the player and to conform to

the intended genre. We ensure that the DED conforms to the intended genre by making the actors responsible for the emergent drama.

For the director and actors' decision mechanism we use Multi-Agent Influence Diagrams (MAIDS) that have been proven to grow linearly with respect to the number of decisions [7]. This is far better than many planning techniques such as STRIPS which is PSPACE-complete [4].

Our first test bed for DED is the classical English murder mystery. The murder mystery genre is a good test bed due to its fairly simple and generic structure and repetitive motifs.

2 Related Work

The novelty of our work is focused around the use of schemas (as we define them) and Multi Agent Influence Diagrams (MAIDS) [7] – which extends Bayesian network and influence diagrams to represent decision problems involving multiple agents.

Our work is in many ways based on the findings of the OZ system [6]. They did a study on live interactive drama with a director, actors, script, observers and player and discovered that it is important to have the player in control of the speed at which the drama proceeds. The player seems to be ready to justify most actions by the characters as being part of the drama, but considers interactions with the director as an intervention.

Facade is comprised of a Drama Manager, beats, characters, story values, actions and natural language processing [10,11]. Contrary to the schemas in DED the beats are explicitly pre-authored in such a way that all actions within the beat are authored explicitly, and the actions for all roles are strictly coordinated to allow for multi-agent coordination ([10], p.45). Additionally all higher level goals and behaviours that drive a character are located in the beats rather than the character. Still the character retains some autonomy when it comes to base level goals and actions such as facial expressions or personality moves ([10], p.45).

The Interactive Drama Architecture (IDA) [9] uses a director and a set characters that enact a fully structured story that is authored by a human author. The director in IDA gives direct commands to the characters, and the characters have very little autonomy.

IDtension [14] bases its approach on narratology and structuralism. The system is authored by defining and scripting a set of tasks that need to be completed in a causal order to complete a certain goal. There are several such causal pathways to complete each goal and thus it provides an emergent drama to a certain degree. It does not use centralized command as DED does, or anything similar to schemas.

GADIN [3] is an emergent system for creating clichéd dramas, as are commonly found in soaps. The system uses central planning to generate actions for characters in the game world with respect to the characters' personalities and current interests, rather than distributing the planning between the characters. The system periodically generates dilemmas to create conflicts in the unfolding drama.

The FATiMA agent architecture [2] is a character based system that uses lessons from live action role-play games (RPGs) and pen and paper RPGs. They emphasises on the need of a Games Master (GM) to guide the emergent narrative without constricting the interactive play. The GM uses hierarchical planning and way-points to help it in plotting the narrative in response to the interactive play between characters. DED uses schemas rather than way-points to structure the drama and distributes the computation between the autonomous decision mechanisms of the actors.

3 The Murder Mystery Drama

The DED architecture is intended for any type of drama, we choose the murder mystery as an initial set-up because it is very structured and contains well known motifs. A typical English murder mystery can be divided into 3 acts, a prologue, a large middle part and an epilogue [15] and it is shaped into a dramatic arc with exposition, complication, climax, fall, and closure [12].

In act I, the exposition, the characters are introduced, the scene is displayed including any secret drawers, hidden compartments, etc. The inciting event is the discovery of the body.

In act II, the complication, the detective interviews all suspects and observes all clues. Many of which are irrelevant to solving the mystery. This means that all clues are revealed by the end of act II.

In act III, climax, fall, and closure, the murderer is revealed by showing that only the murderer had motive, means and opportunity.

We can see from this that there are certain definable goals that need to be fulfilled before progressing from act I to act II and from act II to act III. The three acts can be further divided into even smaller sections, or schemas. Schemas have a much quicker rise and climax than the main drama, albeit less intensity, and thus serve to keep the player engaged [5].

4 The External Structure of Schemas

The director overlooks the emergence of the drama and uses schemas to direct the drama by giving the actors appropriate schemas to play out.

The drama can not move between acts until the objectives of the acts have been adequately satisfied. As an example, the drama will not move from act I to act II until characters' key characteristics have been exposed. If a character is to be intelligent, playful and curious then she needs to have played out actions that are intelligent for a value above a given threshold \mathcal{T} , and the same for playfulness and curiosity.

Skilfully winning a chess game or making 3-4 correct estimates about, for instance, the age of old furniture or showing good arithmetic skills would be sufficient. The algorithm summarises the percentiles to see if it has reached the threshold \mathcal{T} . The director's role is to give the actors an opportunity to show their characteristics by choosing schemas that would be a good fit.

In the example of the murder mystery then characteristics are revealed, but not in direct connection to the true motive behind the character's actions. The player will need to connect motives with actions and characteristics of the suspects. The same applies to revealing clues. The clues should be clearly observable but the player needs to understand how they can be a piece in the puzzle and how to put the pieces together to make a whole picture.

Each schema has a finite set of roles that are annotated as being essential or non-essential. It is only necessary to fill the essential roles to successfully execute a schema, the non-essential roles add variety and increase flexibility. Each role is annotated with a finite set of characteristics that it supports. The characteristics also have a numerical value attached to them, this represents to what degree the display of this characteristic is supported by the role, (see Example II).

The director uses the set of characteristics to match the roles to actors trying to deploy the schemas that best compliment the various characteristics of the characters. The director is not in a good position to make decisions about direct interactions with the player, because the director would need to be constantly aware of everything that takes place – including the internal state of every character in the drama. This would quite rapidly escalate into an intractable computation problem for the director.

The schemas themselves are annotated with the type of clue they can reveal and which act they belong to. The director uses these annotations to filter out schemas that are not appropriate for the current part of the drama. For instance, the director will not consider deploying the *find the body* schema until the victim has been properly introduced.

Example 1 (Interrogation schema). There must be a suspect to interrogate and there can be some witnesses and some policemen present, This schema would be deployed multiple times during the drama. Only the *suspect* is a necessary role besides the player who is the interrogator:

- Drama annotation: can be used in act II only, reveals motive and opportunity to a large degree and means to a small degree.
- Roles: at least one suspect; zero or more witnesses; zero or more policemen.
- Suspect: $\{0.8 * intelligence, 0.7 * gullibility, 0.9 * arrogance, 0.3 * playfulness, 0.1 * competitiveness\}$.
- Witness: $\{0.4 * intelligence, 0.8 * gullibility, 0.9 * arrogance\}$.
- Police: $\{0.8 * intelligence, 0.8 * gullibility, 0.9 * arrogance, 0.9 * competitiveness\}$.

5 The Internal Structure of Schemas

The schemas do not occur in a strict sequential order. Instead the schemas are overlapping and actors typically play out multiple schemas at any given time. The actor uses the schemas to determine which action to carry out. However note that the actions are influenced by the characteristics and state of the character,

which are not imposed by the director or the schema. In this way it can be ensured that actions are always coherent and consistent, even when multiple schemas are active.

The role that an actor plays in a schema is represented by a finite set of actions that the actor can choose from. Some of these actions are essential in that she must play them to complete the schema, and some are not. These actions are annotated by a set of feelings and a set of characteristics that the action can represent. The actors march these sets with their current state of feelings and with their characteristics.

We see an angry woman as *aggressive* if she shakes someone and we see a happy woman as playful if she shakes someone. The action of shaking another character is annotated with both aggressive and playful. The player can be safely trusted to interpret it as intended because they will have seen a happy character laughing, smiling and generally acting in a happy manner, while the angry character will have demonstrated angry behaviour. It may be that a character should show playfulness when happy but not when she is angry. This is achieved with a Bayesian net that the actor uses to gauge the current status of the character she is playing, i.e. the character will be more playful when happy.

The actors' goals are to engage the player and to support the drama progression. The actor is aiming to demonstrate those characteristics of the character they are playing which have an associated value above a given threshold and to meet any other required goals within the unfolding drama. Because the actor is responding directly to the player, she may find that demonstrating characteristics other than those intended by the director will better serve these goals. This is good as it increases the probability that the goals of the act will be reached in response to the actual interactions of the player and other actors.

The schema also contains any knowledge that is necessary or useful in playing the schema in the form of a Bayesian network. The actor uses a Bayesian network as a knowledge base and to describe her beliefs about: other characters; the player; and the state of the drama environment. For more details on this application of Bayesian networks see [1].

6 Outlook and Conclusions

Implementation of these techniques has begun in Libsecondlife [8] and C#, which is a library to program bots that can log in and act in every way as an avatar in the Second Life virtual world [13]. Second Life is a good test bed as it allows access to a very large pool of testers. Libsecondlife has been tested to the extent that we have been able to determine that it is robust enough to support our implementation of DED.

The work completed thus far shows great promise. The schemas are an excellent way of dividing the drama into manageable narrative structures; and defining the roles, action, and knowledge base of the actors.

Complexity is clearly reduced by our architecture as the autonomous agents need only tackle problems that directly affect their own goals rather than

attempting to optimise plans for multiple other agents. Additionally, this type of distributed computation greatly simplifies computing by reducing the size of each problem with clearly defined structures and filters.

References

1. Arinbjarnar, M.: Rational dialog in interactive games. In: Proceedings of AAAI Fall Symposium on Intelligent Narrative Technologies, Westin Arlington Gateway, Arlington, Virginia (2007)
2. Aylett, R., Louchart, S., Tychsen, A., Hitchens, M., Figueiredo, R., Mata, C.D.: Managing emergent character-based narrative. In: The Second International Conference on Intelligent Technologies for Interactive Entertainment, Cancun, Mexico (January 2008)
3. Barber, H., Kudenko, D.: Dynamic Generation of Dilemma-based Interactive Narratives. In: Proceedings of the Artificial Intelligence and Interactive Digital Entertainment conference (AIIDE) (2007)
4. Bylander, T.: The computational complexity of propositional STRIPS planning. *Artificial Intelligence* 69(1-2), 165–204 (1994)
5. Esslin, M.: *An Anatomy of Drama*. Sphere Books, London (1976)
6. Kelso, M.T., Weyhrauch, P., Bates, J.: Dramatic presence. *PRESENCE: The Journal of Teleoperators and Virtual Environments* 2(1) (1992)
7. Koller, D., Milch, B.: Multi-Agent Influence Diagrams for Representing and Solving Games. *Games and Economic Behavior* 45(1), 181–221 (2003) (Full version of paper in IJCAI 2003)
8. Libsecondlife (June 2008), http://www.libsecondlife.org/wiki/Main_Page
9. Magerko, B.: Story representation and interactive drama. In: Proceedings of the Artificial Intelligence and Interactive Digital Entertainment conference (AIIDE) (2005)
10. Mateas, M.: *Interactive Drama, Art, and Artificial Intelligence*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, Technical Report CMU-CS-02-206 (December 2002)
11. Mateas, M., Stern, A.: Build it to understand it: Ludology meets narratology in game design space. In: Proceedings of the Digital Interactive Games Research Association Conference, Vancouver B.C (June 2005) (Included in the Selected Papers volume)
12. Prince, G.: *A Dictionary of Narratology*. University of Nebraska Press (2003)
13. Second life (June 2008), <http://secondlife.com/>
14. Szilas, N.: Stepping into the interactive drama. In: Göbel, S., Spierling, U., Hoffmann, A., Jurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) *TIDSE 2004. LNCS*, vol. 3105, pp. 14–25. Springer, Heidelberg (2004)
15. Todorov, T.: *The Typology of Detective Fiction*. In: *The Poetics of Prose*, pp. 42–53. Basil Blackwell, Oxford (1977)

Developing a Drama Management Architecture for Interactive Fiction Games

Santiago Ontañón, Abhishek Jain, Manish Mehta, and Ashwin Ram

Cognitive Computing Lab (CCL)
College of Computing, Georgia Institute of Technology
Atlanta, GA 30332/0280
{santi,abhishek,mehama1,ashwin}@cc.gatech.edu

Abstract. A growing research community is working towards employing drama management components in interactive story-based games. These components gently guide the story towards a narrative arc that improves the player’s experience. In this paper we present our Drama Management architecture for real-time interactive story games that has been connected to a real graphical interactive story based on the Anchorhead game. We also report on the natural language understanding system that has been incorporated in the system and report on a user study with an implementation of our DM architecture.

1 Introduction

There has been a growing interest in creating story-based interactive fiction games where the player is considered an active participant in the ongoing narratives. The component in charge of guiding the complete dramatic experience is called Drama Manager (DM) [2] or Director [4]. The DM employs a set of actions provided at appropriate points in the ongoing game whereby the player is guided towards certain aspects of the story. Previous approaches to drama management have either not been connected to a concrete world [11] or have been evaluated without using real human players interacting with a real game [7]. In our previous work [9] we evaluated search based DM techniques in a simple implementation of the game Anchorhead [3]. Anchorhead, created by Michael S. Gentry, is a game with a complicated story and several subplots, making it amenable for drama management studies.

Previously, we have shown that player modeling is a key component for the success of drama management based approaches [8] and DM should take into account player’s previous gaming experience and player experience with the current game in its use of strategies. We have also reported [9] that the player model should not only be built using the feedback on interestingness of intermediate game events but also on the player feedback on various strategies used by DM that were visible to the player during the interaction. In this paper we will focus on three main goals. First, we tackle the problem of using search-based DM techniques in real-time games. The vast majority of commercial interactive stories are real-time, so it is very important to take this issue into account. Second,




Fig. 1. A screenshot of our graphical implementation of Anchorhead

most search-based DM techniques assume a limited set of actions the player can execute (in order to cut down the size of the search space that the DM has to explore), in this paper we will explore how to apply DM techniques to a game where the number of actions the player can execute is increased as a natural language interface (in our game) creates a more open ended input action space. Finally, we report a user study with an implementation of our DM architecture and an initial analysis of the study. This report is useful in two ways: a) it supports that our DM architecture is effective in real-time games, with open ended user actions, and b) provides guide for future development of our research.

The rest of the paper we will first present a brief introduction to the game we have used as our testbed, Anchorhead and present our interactive stories architecture. After that, we introduce our technical approaches to NLU, player modeling and drama management, followed by the results from the player evaluations. Finally, we conclude the paper with some final thoughts and future directions.

2 Interactive Stories Architecture

Anchorhead is an interactive story game created by Michael S. Gentry [3]. We have developed a subset of the complete game and it includes a graphical interface for interaction with the player. Graphical as well as text descriptions of the current scenario are presented to the player, who then enters commands in textual format, e.g. “enter the mansion” or “take the key”. For this paper we have focused on a subpart of the story, identified by [7] as interesting for evaluating drama management approaches. Figure 1 shows a screenshot of our Anchorhead implementation. To evaluate our approach we developed a generic interactive




Fig. 2. Overview of the main modules and their relations in our interactive story architecture

stories architecture and implemented a graphical version of Anchorhead with it. Our architecture consists of five modules (shown in Figure 2), namely:

- Graphical Interface (GUI): through which the user interacts with the system, shows a graphical representation of the game (see Figure 1), and allows the player to enter commands in English.
- Natural Language Understanding (NLU): parses the English text and generates a representation that can be understood by the game engine.
- Game Engine (GE): responsible for running the game, maintaining the physical state, story state and a history of what happened during the game.
- Player Modeling (PMM): develops a player model using case-based reasoning techniques from the player actions.
- Drama Manager (DMM): takes the player model and the current game state and generates drama manager actions (DM actions) in order to influence the course of the game towards more interesting plots for the current player.

All the modules in our architecture are developed independently and are domain independent. In order to define a game, the author has to specify (using XML files) the initial state, graphics, the set of plot points and the set of DM actions available.

A game executes in the following way. When the system is launched, the player specifies his previous playing experience (in a 5 point scale) and the number of previous times he has played the game (this is used as part of the player model). Then the player starts playing the game. Each time the player enters an action, the current player trace is updated and the player modeling module updates the player model. The drama manager runs in parallel with the game and constantly is performing time-bounded searches for appropriate DM actions to execute. Each time the DM finishes a search, it executes the resulting DM action, and each time the player performs an action, the DM checks whether the action invalidates its search process or not, and restarts the process appropriately. When the game is over, the player is presented with a feedback form, where he can specify which parts of the game he enjoyed and which parts he didn't. The form allows the player to provide feedback for plotpoints as well as for observable drama manager actions (like hints). The result of the form is stored in the case base of the player modeling module for future games.

2.1 Natural Language Understanding

The text typed by the user is sent to the NLU module where it creates a semantic representation for the input that can be further used by the game engine. The NLU approach is based on techniques used previously in [5,6]. The NLU consists of two main parts: a key phrase spotter and a semantic analyzer. The first stage of processing inside the NLU involves detecting multi-word expressions from a stored set of words labeled with semantic and syntactic tags. After detecting the keyphrases from the user utterance, the processed utterance is sent to the semantic analyzer for further processing. Here, dates, age, and numerals in the user utterance are detected while both the syntactic and semantic categories for single words are retrieved from a lexicon. Relying upon these semantic and syntactic categories, grammar rules are then applied to the utterance to help in performing word sense disambiguation and to create a sequence of semantic and syntactic categories. At the same time, the NLU calculates a representation of the user utterance in terms of dialog acts. Generic rules are defined inside the semantic analyzer for detecting dialog acts. These dialog acts provide a representation of user intent like types of question asked (e.g., asking about a particular place or a particular reason), opinion statements (like positive, negative or generic comments), greetings (opening, closing). The final output is a semantic representation consists of the action that the player wants to conduct in the game and various properties of the actions encoded as attribute-value pairs. Attributes of the actions are things like the location at which the action needs to be conducted, object to which the action should be applied and so on. Table 1 provides two examples of processing inside the NLU.

Table 1. Two Examples of processing inside the NLU

NLU Submodule	Submodule Output
User Input	Could you take me to the living room
Semantic Analyzer	<action_name=go_to> <attribute name=location value=livingroom> <attribute name=dialogact value=question>
User Input	give me the flask please i need it
Semantic Analyzer	<action_name=give> <attribute name=topic value=player> <attribute name=object value=flask>

2.2 Player Modeling Module

The player modeling module (PMM) maintains a player model for the current player based on the feedback provided by players at the end of each game. This feedback contains player opinions on the game, including the parts they enjoyed and those that were not interesting from their perspective. At the end of each interaction, the PMM stores this player feedback along with the corresponding trace of player actions during the game. When a new player interacts with the

game, the stored feedback of previous players is used to predict which storylines and game elements would be enjoyable for the new player.

Specifically, before starting the game, the player enters his overall gaming experience and the number of times he has played the Anchorhead game, and at the end of each game, the player is presented with a sequence of the plot points that he visited over the course of the entire game episode. From the list, the player is asked to select his preference of the plot points based on a 5 point Likert scale: *strongly like*, *like*, *indifferent*, *dislike* and *strongly dislike*. After that, the player is presented with the list of observable DM interventions (mostly hints) that occurred during the game. The player rates them on a 5 points scale as well. Finally, the player rates the game as a whole.

We use a case-based reasoning (CBR) approach [1] for the PMM module. Based on the feedback provided by the player, the systems builds a *player preference model* that models the stories and DM actions that the player is likely to enjoy. CBR works by reusing previous experiences (called cases) to solve new problems. In our CBR approach for player modeling each time a player plays a game and provides feedback, the trace of the game plus the feedback is stored in the form of a *case*. As a particular player is playing the game, his current trace is compared to the traces in the different cases stored in the case base. The PMM retrieves the most similar traces and creates a player model by aggregating the player feedback contained in them. The assumption made here is that the trace of a player captures its playing patterns, and that players with similar playing patterns will have similar preferences. Specifically, cases in the PMM consist of the following elements: a) Player gaming experience, on a 5 point scale; b) number of times the player had played Anchorhead before; c) the game trace, consisting of a sequence with single action that the player executed, and each DM action that the DM executed in the game; and d) all the information provided by the player in the feedback form: overall experience rating, confidence, plot-point preference and DM action preference.

The two key processes in the PMM are: a) how to compare the trace of the current player with the traces stored in the cases in the case base, and b) once a subset of cases have been identified as relevant, how are they used to generate a player model. Let us explain those two processes in detail.

Case Retrieval. Case retrieval works by computing a distance metric among the current player and previous players (stored as cases) and selecting a subset of k cases with the minimum distance. In order to compute such distance metric, the PMM uses 4 different distance metrics, that are averaged in order to compute the final distance: a) difference in player gaming experience, b) difference in the number of times the player has played the game before, c) difference of the player action traces, and d) difference of the DM action traces.

Each of the distance metrics generates a number between 0 and 1. The metrics for gaming experience and number of times the player has played the game are simple, so let us explain how do we compute distances among traces.

To compute how different a player action trace is from another we perform some preprocessing (as shown in Figure 3). Specifically, if we are comparing a

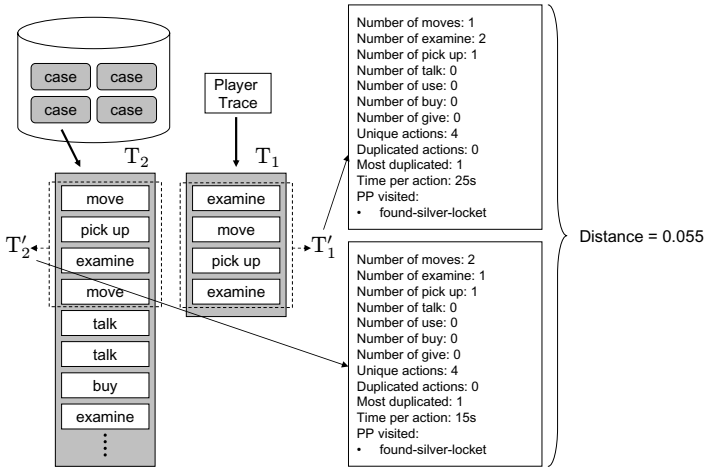


Fig. 3. Computation of player action trace distance

trace T_1 with a trace T_2 , having n_1 and n_2 actions respectively, we only consider the first $\min(n_1, n_2)$ actions of each trace. Let's call these reduced traces T'_1 and T'_2 . Now, T'_1 and T'_2 have the same number of actions and can be properly compared. The rationale behind this process is that when a player has just started to play the game and has performed only a few actions, it does not make sense to compare its trace against the full trace of another player.

From T'_1 and T'_2 , we compute a set of features consisting on: the number of different kinds of actions (movement, talking, using, etc.), number of unique actions, number of duplicated actions, the highest number of times a player has repeated an action, the average time a player takes to decide an action, and the difference among the set of plot points explored in T'_1 and T'_2 . Each of those features define a number, and to compare two traces we simply compute a normalized Euclidean distance, which generates a number between 0 and 1.

In order to compare two DM action traces, we perform the same preprocessing as for player action traces and obtain the reduced DM traces, and define an analogous set of features.

Player Model Generation. The output of the retrieval process is a subset of k cases ($k = 3$ in our experiments) relevant to the current player: c_1, \dots, c_k , and their distances to the current player d_1, \dots, d_k . A player model in our system consists only of three elements: a) a list of predicted plot point interestingness, b) a list of predicted DM action interestingness, and c) a confidence value in the interval $[0, 1]$, corresponding to how sure the PMM is that the current module is correct. To build such model, the interestingness value for each plot point, and DM action is computed by aggregating the interestingness values from cases by using the similarity metrics to weight the individual interestingness values, i.e. the interestingness values in case c_i have d_i as their weight. The confidence is computed in the same way by averaging the confidences in the cases weighting

them by their distances. The output of the PMM is a player model that consists of the predicted interestingness of each plot point for the current player and also a confidence on this player model.

2.3 Drama Management Module

Given the player preference model, the current game state, and optionally some author specified story guidelines, the Drama Management Module (DMM) plans story arcs that maximize both the player model and the author specified story guidelines. The DMM module constantly starts searches in parallel with the game using this information to select, if necessary, a particular action to influence the story towards the story arc identified by the DMM as the most interesting. The game author specifies a set of *drama manager actions* (DM actions). These actions represent the things that the drama manager can carry out to influence the game, e.g. “prevent the player from entering the library by locking the door” or “make the bar keeper start a conversation with the player about the suspicious person”. The DM actions can be classified in several groups:

- *hints*: actions that has no direct effect on the game, but hint the player towards a particular direction.
- *causers*: forces something to happen in the game instead of waiting for the player to do it.
- *deniers*: prevents the player from doing something.
- *temporary deniers*: only prevent the player from doing something temporally.
- *reenablers*: reenables a previously denied line of action.

Notice that some of the actions (such as deniers or causers) have to be subtle so that the player does not feel manipulated.

Each DM action consists of several elements: a name, the type, a set of preconditions, and a set of effects. The preconditions of a DM action are boolean expressions that can test whether some plotpoints have already been visited, if the player is in any particular location, or has a particular object, etc. The effects that a DM action can have are: causing characters in the game to execute actions, firing plotpoints, directly modifying the game state, or hinting actions.

An example of DM action in our implementation is “bum_hints_crypt_key_is_in_basement”, that has as preconditions that the player is in the park, that the player has bribed the bum and that the player does not have the crypt key. The effects of the actions is causing a particular character in the game, the bum, to tell the player that there is a key hidden in the basement. Specifically, this action is a *hint*. As the game is going on, the DM will choose to execute this action if it realizes that by providing the key to the player, it will cause the player to reach enjoyable plot points.

The algorithm used the the DMM is a variation of the algorithm explained in [9], based in an *expectimax* method, but adapted to real-time games in which the player can type actions through natural language.

Searching For an Interesting Story Arc. In order to decide which DM actions to execute, the DMM performs a search in the space of possible stories. Such search is represented as a search tree with alternating DM actions and player actions. The result of the search is an expected degree of interestingness of the different stories that will unfold depending on which DM action the DMM executes. Then, the DMM will select the DM action that maximizes such interestingness. Note that the *no-op* DM action (that represents executing no DM action) is always one of the DM actions considered.

In order to perform this search, the DMM uses a hybrid *montecarlo expectimax* algorithm, illustrated in Figure 4. The expectimax algorithm performs search by alternating two types of nodes in the search tree: *max* nodes and *exp* nodes. In the leaves of the search tree, the DMM will evaluate the interestingness of the story that unfolded to reach that leaf, and the interestingness value will be propagated up. In the *max* nodes, the interestingness is assigned as the maximum of the interestingness of the children nodes, and in the *exp* nodes, it is assigned as the average of the interestingness of the children nodes. Naturally, *exp* nodes will be nodes where the DMM is considering player actions, and *max* nodes correspond to nodes where the DMM is considering DM actions. The first node corresponds to the current game state, and is a *max* node. The next set of nodes is expanded by taking into account each possible DM action that can be performed in the current game state. This generates the first layer of nodes (*ply1*), where all the nodes are *exp* modes, as shown in Figure 4. The second set of nodes is generated by considering all the possible player actions, this generates the second layer (*ply2*), where all nodes are *max* nodes.

Since searching in such search tree has a very high cost, the DMM only performs exhaustive search for a predefined number of layers (set to 3 in our experiments). After that, the DMM performs *montecarlo search*. Montecarlo search estimates the interestingness value by expanding single branches at random, but at a very high depth (set to 20 in our experiments). The number of branches




Fig. 4. Visualization of the search process performed by the DMM, where the DMM had to decide among three DM actions

that the montecarlo part of the DMM expands is variable and depends on the amount of time left for search (as explained in the next section). At the end of each one of the montecarlo branches, the DMM uses the player model and the set of author specified guidelines to compute an interestingness value for that branch. Such interestingness is computed in the following way: first, the DMM uses the player model to predict the interestingness of each plotpoint visited in that branch, and each visible DM action executed in that branch, this gives a value pmi . Second, the DMM uses the author specified guidelines to compute a second interestingness value asi . The final interestingness is computed as $i = c \times pmi + (1 - c) \times asi$, where c is the confidence in the player model. As figure 4 shows, such interestingness values are propagated up the search tree, and finally the DMM will select the DM action that maximizes the interestingness. In the example shown in Figure 4, that corresponds to DM action dma_1 , that achieves an interestingness value of 0.4.

Search-Based Drama Management on Real-Time Games. Previous search-based drama management approaches have been implemented in turn-based games, where each time the player executes an action, the DM has time to perform a search and return a DM action. In our game the DMM is constantly running in parallel with the game as it evolves. In order to support this, the DMM constantly executes time bounded searches (2 seconds per search in our experiments). Each time a search is finished the DMM executes the proposed DMM action (if any). Each time the player (or any character in the game) executes an action, the DMM verifies that this action does not invalidate the search, and in case it is invalidated the search is restarted. In order to use the search time efficiently, the DMM uses an iterative deepening method that searches first only at depth 1, then depth 2, etc. As soon as the maximum level for exhaustive search is reached, then the DMM executes more and more montecarlo samplings, until the time is over.

Another challenge faced is that actions in our game take time, i.e. moving the character from a location to another takes some time (depending on the distance it has to travel), thus the DMM cannot just open a search tree where each node is a cycle of the game. In order to solve this, we incorporated the concept of *stable game state*. A stable game state is one where no character in the game is executing any action. The DMM always starts its search in a stable game state, and when an action is selected in a branch of the search, the game state is simulated until we reach another stable game state, from where the search can continue. Figure 4 illustrates the effect of this. On the left hand side we have the number of game cycles, and we can see that each level of the search tree does not correspond exactly with a particular game cycle.

The last challenge is related to the amount of actions the player can execute. Since in our game the set of actions is practically unbounded (the player can potentially command anything through natural language, restricted only by the limitations imposed by our NLU module and by the set of primitive actions supported by our game engine). Notice for instance, that the player can type actions such as “go to the bar”, “examine the ground”, “say hello to the bum” or

“ask the bum about his past”; just the set of possible communication actions to interact with other characters is unbounded. Since the DMM uses search based techniques (similar to those used in Chess playing programs), it needs a set of actions that the player can execute to perform the search. In our system, the DMM analyzes the set of plotpoints and the current game state to generate the set of actions that can affect the story. The rest of actions are considered as “no-effectors” by the DMM and are not considered in the search. Using this method, the author does not have to specify any list of possible actions (as in most previous DM approaches), the DMM can figure them out by itself.

3 Evaluation and Initial Analysis

We conducted our experiments in two batches. In the first, we invited five of the participants from [9] who had previously played an earlier turn-based text version of the game with an earlier version of our architecture. The second batch of ten participants had never played Anchorhead before.

The players in the first batch were interviewed at the end to get feedback on the usability improvements we made from the earlier version of our system namely shifting to a new graphical interface and employing a natural language based interaction modality. The first users acted as pilot subjects that helped us improve the graphical interface, NLP and DM modules. The suggested improvements from the first batch were implemented before the second batch of participants played the game. One of the improvement from their interaction was to improve the performance of the NLU by adding more lexical entries and rules. During their interaction, the users felt that the language based interface seemed like a more natural way of interaction and was much more interesting than the text based interface (where the user simply selected from a fixed list of possible actions presented to him as a choice list). However, the open ended nature of the NLU made it difficult for the users to figure out appropriate actions. In order to make it easier for users to find right set of actions, we improved the language interface with a help screen that provided them with example commands and their usage. We also added some more DM hints to help the next batch of users at the points where the first batch of users seemed to be generally lost.

After the improvements, we carried out a study with a second batch of participants to measure the effect of DM strategies on player’s experience. For this batch of the study, we recruited 10 participants with a range of genders (2 females and 8 males). Each player was provided with an explanation on Anchorhead and asked to sign a consent form before starting the game. The player filled a background questionnaire to obtain information such as previous gaming experience or types of games they liked to play. During each episode, a researcher logged his observation of player actions and any unusual reactions. Each player played twice, once with DM active and another one without DM. On an average, the complete player interaction (both game playing episodes) lasted for about 45 minutes each. At the end of each episode, the player was asked to provide an interestingness value and an associated confidence value on a 5 point Likert scale

for the overall game experience as well as the intermediate story events that were encountered during the interaction. At the end of both gaming episodes, participants were interviewed about their experience. The evaluation accounted for the order in which players played with and without DM by making half of the subject play first with DM and then without it, and the other half play first without DM and then with it. We transcribed the player responses from the interviews and observed players actions during the game episodes and analyzed them using a well-known qualitative analysis method, Grounded Theory [10] as used in [9]. Next we present results from qualitative analysis. We are planning to conduct a complete quantitative analysis in the future.

Concerning the interaction with the game, players found the natural language interface easy to use, although sometimes it was hard for them to figure out the appropriate actions due to the open ended NLU. Most players preferred the NLU interface to the choice-based interface (although a minority complained that it was annoying to type the commands). Further, having an open ended NLU interface made the players more engaged in the game as successfully taking an action was more rewarding (in the choice based interface it was obvious to figure out which actions to take next). Some players were frustrated as the open-ended NLU interface made them believe that they could interact freely with other characters, and in reality those character had very limited conversational topics (the characters in our game trigger responses or actions if the player talks to them about certain keywords or gives them certain objects).

Concerning the observable DM interventions some players felt that the DM was providing too many hints. Others liked the hints as they guided them in the game (when played without the DM some players were totally lost). Some other players felt that some hints were too blunt and told more than expected. (which points out that the DM should have several degrees of hints from more subtle to more blunt). Most players followed the hints and were able to finish the game when the DM was active, whereas without the DM most of them got stuck. Players felt a positive effect of the system's interventions on their interaction, when they saw that when playing without DM they ended up in an ending they didn't like, whereas playing with DM, the system (DM) prevented that ending from happening resulting in a more enjoyable experience. The DM hints also allowed certain players, who were stuck otherwise to finish the game.

4 Conclusions

In this paper we have presented an improved version of the drama management architecture initially presented in [9]. Specifically, our improvements include: real-time drama management, improved user modeling including preferences about DM actions, natural language interface for open ended user interaction (improving the DM to support this), and a graphical representation that is more appealing to non-gamers. We performed a user evaluation and initial qualitative analysis with 15 users (5 users who had used the old system, and 10 new users).

The main conclusion we can draw from our experiments is that DM techniques are applicable to real-time games, and that DM generally improves player experience. Moreover, our user evaluation reinforced the fact that the DM should take into account player experience when deciding how to influence the story and also when deciding which hints to provide to the player.

As future work, we are in the process of conducting a quantitative analysis of our new DM architecture to measure how much the player experience improves with and without DM, and also how much it is due to the improved DM architecture with respect to the one reported in [9]. We are also in the process of developing a bigger game based on our architecture and plan to make it available on-line in order to have a larger number of players and perform further and better evaluations of the system.

References

1. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. *Artificial Intelligence Communications* 7(1), 39–59 (1994)
2. Bates, J.: Virtual reality art and entertainment. *The Journal of Teleoperators and Virtual Environments* 2(1), 133–138 (1992)
3. Gentry, M.S.: Anchorhead (1998), <http://www.wurb.com/if/game/17.html>
4. Magerko, B., Laird, J., Assanie, M., Kerfoot, A., Stokes, D.: AI characters and directors for interactive computer games. In: *Proceedings of the 2004 Innovative Applications of Artificial Intelligence Conference* (2004)
5. Mehta, M., Corradini, A.: Understanding spoken language of children interacting with an embodied conversational character. In: *Proceedings of the Combined Workshop on Language-Enabled Educational Technology and Development and Evaluation of Robust Spoken Dialog Systems at ECAI 2006*, pp. 51–58 (2006)
6. Mehta, M., Corradini, A.: Developing a conversational agent using ontologies. In: Jacko, J.A. (ed.) *HCI 2007. LNCS, vol. 4552*, pp. 154–164. Springer, Heidelberg (2007)
7. Nelson, M., Mateas, M., Roberts, D., Isbell, C.: Declarative optimization-based drama management in interactive fiction. *IEEE Computer Graphics and Applications* 26(3), 33–41 (2006)
8. Sharma, M., Mehta, M., Ontañón, S., Ram, A.: Player modeling evaluation for interactive fiction. In: *AIIDE 2007, Workshop on Optimizing Player Satisfaction*. AAAI Press, Menlo Park (2007)
9. Sharma, M., Ontañón, S., Mehta, M., Ram, A.: Drama management evaluation for interactive fiction games. In: *AAAI 2007 Spring Symposium on Intelligent Narrative Technologies*. AAAI Press, Menlo Park (2007)
10. Strauss, A., Corbin, J.: *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*. Sage, Thousand Oaks (1990)
11. Weyhrauch, P.: *Guiding Interactive Drama*. PhD thesis, Carnegie Mellon University (1997)

Planning and Interaction Levels for TV Storytelling

Angelo E.M. Ciarlini¹, Marcelo M. Camanho¹, Thiago R. Dória¹,
Antonio L. Furtado², Cesar T. Pozzer³, and Bruno Feijó²

¹ UniRio – Departamento de Informática Aplicada
Av Pasteur, 458 – Térreo – Rio de Janeiro, Brazil
{angelo.ciarlini,marcelo.camanho,thiago.doria}@uniriotec.br

² PUC-Rio – Departamento de Informática
Rua Marquês de São Vicente, 225 – Rio de Janeiro, Brazil
{furtado,bfeijo}@inf.puc-rio.br

³ Universidade Federal de Santa Maria - Dep. de Eletrônica e Computação
Campus Universitário – pr. 07, 3º and. Santa Maria, RS, Brazil
pozzer@inf.ufsm.br

Abstract. Interactivity, coherence and diversity of the stories are key issues for the development of interactive storytelling systems. When stories are to be told via interactive TV, special interaction methods are also necessary in order to cope with high responsiveness requirements. In this paper, we describe the extension of the interactive storytelling system Logtell we have developed to run in an interactive TV environment. Planning algorithms have been applied to provide coherence and diversity for the stories at levels of both plot composition and dramatization. A new architecture was designed to combine these algorithms with special interaction methods incorporated to achieve high responsiveness.

Keywords: Storytelling, Planning, Interactive TV, Logics.

1 Introduction

Different approaches have been proposed for interactive storytelling, some of them closer to games and others to filmmaking and Literature. In *character-based* storytelling systems [1], the story emerges from the real time interaction between autonomous agents, each one with its own objectives. The main advantage is to facilitate user intervention, since any actions of any character in the story can be influenced, so that the plot may take different directions. In these systems, stories are usually told with a first-person viewpoint and users play the role of characters, as in games. The main challenge for this approach is keeping the coherence of the stories, since user intervention might create inconsistencies. Some other systems have a *plot-based* approach [2,3], wherein a series of rigid rules, built into the plot, guide the narrative, making user intervention far more limited. In such approach, there is a stronger control over the story being presented, preventing the user to stray from the context defined by the author. The general given structure usually includes beginning, middle and ending points previously fixed by the author, and user interaction affects only the

way the story will reach these predefined points. One of the main inspirations for this kind of model has been the seminal literary work of Vladimir Propp [4], at the beginning of the twentieth century. Propp examined a large number of Russian fairy tales, and showed that they could all be described by specializations of 31 typical *functions*, such as villainy, hero's departure, reward, etc. In purely plot-based models, the intervention of the user is more limited, but it is much easier to guarantee coherence together with a measure of dramatic power. There are also alternatives that integrate characteristics of both plot-based and character-based approaches. The interactive system *Façade* [5], for instance, has a drama manager that keeps the characters largely autonomous most of the time, but changes their behaviour to move the plot forward, conciliating higher-level goals, which are essential to the story, with lower-level goals, specific to the autonomous behaviour of the characters. *Erasmatron* [6], in contrast, starts from the notion of verbs and sentences. Actions are represented by verbs with roles assigned to characters to form sentences.

Automated planning algorithms are important parts of some storytelling systems, because they can be used to create a logical chaining of events. In [1], for instance, hierarchical task network (HTN) planning is used to control the way characters achieve their goals in accordance with user intervention. HTN planning tends to be efficient but less general, demanding the previous construction of a task hierarchy and methods to perform each task. More flexible planning algorithms have also been adopted, as in [7] for example. Such algorithms do not limit solutions to previously defined alternatives; instead of that, they combine events conciliating different objectives with pre-conditions and effects of each event. The computational effort, however, becomes even greater, due to the complex nature of automated planning.

Logtell [8] is an interactive storytelling system we have developed to create plots in accordance with a formal description of the genre and the initial situation of the story. *Logtell* is based on modelling and simulation. Its basic idea is to capture the logics of a genre and then verify what kind of stories can be generated and told by simulation combined with user intervention. Specifically, we try to conciliate both plot-based and character-based modelling. On the one hand, we borrowed from Propp's ideas, but tried to extend his rather informal notion of function. Typical events are described by parameterized operations with pre-conditions and post-conditions, so that planning algorithms can be used for simulation. On the other hand, character-based modelling is added under the form of goal-inference rules for each kind of character, specifying how situations can bring about new goals for the characters. In the first version of *Logtell*, our concern on diversity of stories was focused on plot composition. When plots were totally or partially generated, they could be dramatized via an animation of virtual actors in a 3D scenario, but diversity and interaction in the dramatization phase was limited.

The second version of *Logtell* has been implemented to run in an interactive TV environment, complying with requirements of high responsiveness, diversity in the dramatization of stories (not only in their composition) and new interactivity options. A new architecture was proposed, in which *Logtell* resorts to automated planning at various levels to create plots, to control the dramatization of scenes and to animate virtual actors in a 3D scenario. By using plans based on formal specifications, we try to guarantee coherence and diversity for plot composition and dramatization. At the same time, we try to offer the users the opportunity to interact at various levels in both active and passive fashions, which might be needed when the user just wants to watch the story unfold by itself.

In this paper, we present the levels of automated planning and their relation with interactivity in the new version of Logtell. Section 2 presents new Logtell's overall architecture. Section 3 describes our use of automated planning for plot composition. Section 4 describes the interaction methods for plot composition in the new architecture. Section 5 shows how dramatization has been structured to provide more diversity for stories and additional opportunities for user interaction. Section 6 contains our concluding remarks.

2 New Logtell's Architecture

Logtell adopts a client-server model as shown in Figure 1. The client-side is responsible for user interaction and dramatization of stories. The application server side is responsible for simulation. All processes run in parallel and are coordinated with each other. For each story being told there is an application running in the server, while one or many applications are kept running in different clients. This takes care of the case wherein multiple users are simultaneously sharing the same story. If clients are set-top boxes for interactive TV, their computational resources are limited, making it hard to perform CPU-intensive tasks such as automated planning. By concentrating simulation tasks in application servers, it is easier to achieve higher scalability. In addition, it is also easier to exert control when a single story is shared by many users.

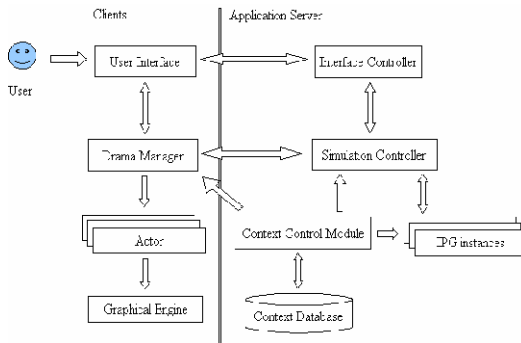


Fig. 1. Logtell's Architecture

The access of all modules to the context of the stories, specified in the Context Database, is always performed via the Context Control Module (CCM), which runs in the server. The context contains the description of the genre, according to which stories are to be generated, and the initial state, specifying characters and the environment at the beginning of the story. The genre is basically described by: (a) a set of parameterized basic operations, with pre- and post-conditions, logically specifying the events that might occur; (b) a set of goal-inference rules, specified with a temporal modal logic, which define situations that lead characters to pursue the achievement of goals; (c) a library of typical plans, corresponding to typical combinations of operations for the achievement of specific goals, which is organized according to "part-of"

and “is-a” hierarchies; (d) logical descriptions of initial situations for the stories, introducing characters and their initial state; (e) a nondeterministic automaton for each operation specifying alternative ways events based on it can be dramatized; and (f) graphical models of 3D virtual actors.

Plot generation is performed by the Interactive Plot Generator (IPG). IPG generates plots as a sequence of chapters. Each chapter corresponds to a cycle in which user interference is incorporated, goals are inferred and planning is used to achieve the goals. IPG is controlled by the Simulation controller. Multiple stages, each one corresponding to a chapter, usually occur in order to generate a plot.

On the client side, the user interacts with the system via the User Interface, which informs the desired interactions to the Interface Controller placed at the server side. The Drama Manager requests the next event to be dramatized from the Simulation Controller, and controls actor instances for each character in a 3D environment running on our Graphical Engine. On the server side, the Interface Controller centralizes suggestions made by the various clients. When multiple users share the same story, interactions are selected according to the number of clients that requested them. When there is only one client, suggestions are automatically sent to the Simulation Controller. The Simulation Controller is responsible for: (a) informing the Drama Manager, at the client side, the next events to be dramatized; (b) receiving interaction requests and incorporating them in the story; (c) selecting viable and hopefully interesting strong interactions to be suggested to the users; (d) controlling a number of instances of the Interactive Plot Generator, in order to obtain the next events to be dramatized; and (e) controlling the time spent during simulation.

In the new architecture, there can be various instances of IPG running on the server. Besides the instance corresponding to the current flow of the story, others are used to avoid interruption in the dramatization. The simulation has to be some cycles ahead of the dramatization to keep responsiveness. When there is no user intervention, goals are inferred and events are planned continuously. When users interact with the system, however, they interact in accordance with the events that are currently being dramatized. The Simulation Controller keeps snapshots of the state of the simulation after each cycle, so that simulation can be resumed from the correct point after an intervention. Logical coherence of a requested intervention is always checked before being incorporated, or either discarded if inconsistent. When an intervention is incorporated, simulation has to discard simulation cycles that were previously planned without taking the intervention into account. In order to be prepared for interventions, the system creates other instances of IPG, simulating the incorporation of strong interventions to be suggested to the users. But the suggestions are only communicated if the IPG instance confirms that they are consistent. And if they are accepted, the next events are already planned and there is little risk of interruption.

The time spent for simulation is constantly monitored by the Simulation Controller. When there is risk of interruption in the dramatization because there are not enough events planned, a message is sent to the Drama Manager, so that strategies are used to extend the dramatization of the current events until the situation is normalized.

3 Planning for Plot Composition

IPG semi-automatically generates plots of narratives of a certain genre as a simulation process. The initial situation chosen in the context database is the starting point for the simulation. Facts from the initial configuration might be modified by *events* (denoted by instances of operations). The library of operations specifies the kinds of events that may occur in the narratives, designed in anticipation of the character's goals. For each class of characters, there are goal-inference rules, specifying, in a temporal modal logic formalism [9], the goals that the characters of the class will have when certain situations occur during a narrative. It is important to notice that the rules do not determine the specific reaction of a character. They only indicate goals to be pursued somehow. The events that will eventually achieve the goals are determined by the planning algorithm.

The generation of a plot starts by inferring goals of characters from the initial situation. Given this initial input, the system uses a planner that inserts events in the plot in order to allow the characters to try to fulfill their goals. When the planner detects that all goals have been either achieved or abandoned, the first chapter of the story is almost finished. During the generation phase, plots are represented by partially-ordered sets of events. Before dramatizing, however, the total order of the events has to be determined, as explained in Section 4. If the user does not like the story, IPG can be asked to generate another alternative for a chapter and to develop the story from this point on. If the user does not interfere in the process, chapters are continuously generated by inferring new goals from the situations generated in the previous chapter. If new goals are inferred, the planner is activated again to fulfill them. The process alternates goal-inference and planning until the moment the user decides to stop or no new goal is inferred. Users can also interfere in the process by choosing alternatives and forcing the occurrence of events and situations as described in Section 4. Notice that, in this process, we mix forward and backward reasoning. In the goal-inference phase, we adopt forward reasoning: situations in the past generate goals to be fulfilled in the future. In the planning phase, an event inserted in the plot for the achievement of a goal might have unsatisfied pre-conditions, handled through backward reasoning. To establish them, the planner might insert previous events with further unfulfilled pre-conditions, and so on.

We use a non-linear planner that works in the space of plans, not assuming a total order for the events. It was implemented in Prolog, adapted from [10], with extensions. A non-linear planner seems more suitable because it uses a least-commitment strategy. Constraints (including the order of events) are established only when necessary, making easier the conciliation of various goals. Features to permit the abandonment of goals were included, and also constraint programming techniques for dealing with numerical pre-conditions. The planner performs a heuristic search for good plans. It works on many plans in parallel and, at each time, selects the candidate with minimal estimated cost for achieving a complete solution, i.e. a configuration in which all goals and all preconditions of all events are satisfied. It then selects a pre-condition of an event that is not necessarily true and tries to make it true. While doing that, the planner generates all possible successors of the current plan. It considers the possibility of using events already in the plan to establish a pre-condition as well as

the insertion of new events. All possibilities for solving conflicts between events in the establishment of pre-conditions are considered.

To meet the requirement of high responsiveness, efficiency of the planning algorithm becomes essential, and possible enhancements to the planner deserve special attention. Most of the successful cases in real-world applications of automated planning are based on algorithms that use hierarchical task networks (HTNs), such as [11, 12]. These algorithms tend to be more efficient because they reduce the search space. They depend however on the previous definition of a hierarchy of tasks and methods to perform the tasks. The task network has to be built for each domain, and all generated solutions are limited to combinations of previously defined ones, giving less flexibility to the creation of alternatives. Algorithms based on HTNs can work with a partial order of events, and HTNs are compatible with the way we specify a hierarchy of typical plans, used for plan-recognition as explained in Section 4. Taking advantage of that, we are enhancing our planner, producing a hybrid planner capable of working with our type of simulation, mixing non-linear planning with HTN planning.

Typical plans are abstract events usually corresponding to various alternatives of combining basic events to achieve ordinary goals. They can have pre-conditions and effects, but these pre-conditions and effects differ from those of basic events because they are necessary instead of sufficient, that is, additional pre-conditions and effects can be observed in a specific alternative.

When a goal has to be achieved, the system might choose either a basic event or an abstract one to achieve it. In this way, abstract events are initially inserted as if they were basic. When a plan is complete but contains abstract events, HTN planning is applied to create alternatives specializing each abstract event into a set of compatible partially-ordered basic events. Heuristics are used to determine whether basic or abstract events should be chosen for the achievement of a goal (or subgoal). Abstract events tend to be preferred, because a single abstract event is usually necessary to achieve a specific goal. In this case, the planning process is essentially reduced to HTN planning. By doing this, performance is enhanced and, at the same time, we keep the generality and flexibility of our original simulation tool.

In recent years, some neoclassical planners that work in the space of states, with planning-graph techniques [13,14], have given good results. The use of heuristics and control strategies to prune the search space, as in [15,16], are also promising strategies. We are also studying the possible adoption of these techniques.

4 Interaction Methods for Plot Composition

The underlying philosophy of the system consists of providing the user with efficient means for exploring coherent alternatives that the story may allow, and for guiding the plot at the level of events and characters' goals. Interaction can occur in a step-by-step mode, in which the user can interfere after the generation of each chapter, or in parallel with dramatization, in continuous mode.

Step-by-step interaction is analogous to debugging a program. When it is activated, the graphical interface presents a graph with the plot generated after each chapter. Figure 2.(a) shows the graphical interface for this kind of interaction. Each event is represented by a rectangular box that may assume a specific color according to its

current status. As explained in the previous section, plots result from goals that the characters aim to achieve. At each simulation step (i.e. a chapter), new goals may be inferred and automatically added to the plot, which causes the insertion of a new set of events. In step-by-step mode, the events inserted in the plot so far are sent to the graphical interface for user intervention, which offers two commands for automatic plot generation: *another* and *continue*. Command *another* requests from IPG an alternative solution to achieve the same goals of the chapter just finished. Command *continue* asks IPG to try to infer new goals and continue the simulation process. This weak form of intervention usually leads the plot to situations that the author of the context has devised beforehand. The Plot Manager also offers two complementary means for strong intervention in the creation of more personalized stories. Commands *insert situation* and *insert event* allow users to specify situations (specified as goals to be achieved) and events that should occur at specific times. The specific details of how the goal will be accomplished or the event is inserted are left to IPG, which might insert further events to obtain a consistent plot. Insertions are rejected when they are inconsistent or search limits currently configured in IPG are exceeded. In step-by-step mode, the user has to convert the partially-ordered generated plan into a strict sequence, that can be dramatized. To determine the sequence, the user connects the events in a sequential order of his/her choice, respecting the temporal constraints supplied by IPG. Once the current plot (or part of it) is thus connected into a linear sequence, it can be dramatized by invoking the Drama Manager with the *render* command.

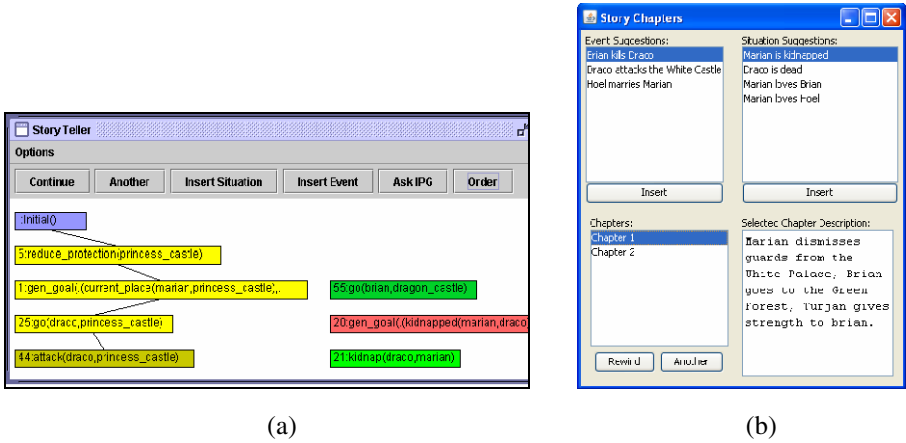


Fig. 2. Interfaces for interaction: step-by-step mode (a) and continuous mode (b)

Considering that our storytelling system is designed to run on an environment such as interactive television, viable alternatives of interaction should not hinder the experience of watching the dramatization. Interaction cannot demand a high level of attention, as in step-by-step mode, unless the user opts to halt the dramatization in order to interact. The interruption of dramatization to allow the user to interact, which was mandatory in the first version of Logtell, is still allowed, but it is expected to be an exceptional case, being replaced by other more expedient kinds of interaction in

continuous mode. The Simulation Controller plays an important role in the implementation of the new forms of weak and strong interactions.

Weak interactions work basically around the “normal” flow of the story, as one would have with *continue* and *another* commands in the first version of Logtell. The Simulation Controller directs the flow of the story by automatically selecting alternatives and total order of the events to be dramatized. Such selections can be done either in a random way or based on user satisfaction models. The idea is that stories are worth of telling even if the user only watches the dramatization, with no intervention. Even in this case, we keep the possibility of watching different but still coherent stories based on the same initial configuration.

Other kind of weak interaction is to return to a previous chapter in the narrative, so that alternative directions for the plot can be chosen. In step-by-step mode, at the end of a simulation phase, the user could examine the new events not yet incorporated, and decide whether or not to consider them interesting; if not, the user would ask for the generation of an alternative. In continuous mode, with parallel dramatization, it becomes highly desirable to extend the backtracking range, to allow the user to undo the narrative up to any previous stage, and have a chance to find how it would develop if different alternatives were chosen at such point. For this objective, different snapshots of the simulation process are kept in memory by the Simulation Controller, corresponding to the end of each chapter.

Suggestions of strong interactions correspond either to the insertion of specific events in the plot or to the directive that a certain situation should occur at a specific time. These suggestions can be made based on the events already inserted in the story and on an analysis of the context of the genre. We shall describe two methods for their creation. Independently of the method utilized, the Simulation Controller checks the consistency of the suggestion with the current narrative, before sending it to the user. The first method to obtain a suggestion of strong interaction uses our library of typical plans. Typical plans usually consist of certain combinations of events whereby the various characters pursue their goals, but they can also correspond to motifs, i.e. recurring structures compiled in the course of critical studies on the genre [17]. IPG contains a procedure for the recognition of plans, based on an algorithm specified by Kautz [18]. The procedure is able to discover that some given events are compatible with a motif for which we have a typical plan, enabling the Simulation Controller to suggest the inclusion of additional events contained in the plan. The second method to obtain a suggestion of strong interaction is based on the application of the goal-inference rules against the current plot. A suggestion, in this case, is triggered by observing that the insertion of a fact at a specific time might lead to the inference of a new goal.

Figure 2 (b) shows the interface used for user interaction in continuous mode. Events and situations to be incorporated in the next chapter are continuously suggested to the user. By selecting a suggestion, the user asks the system to try to force a certain event or situation in the next chapter. If it is logically possible, the new suggestion is likely to be incorporated as a strong interaction; if not, it is discarded. By selecting a chapter and pressing button *rewind*, the story backtracks up to the desired point. By pressing *another*, the story also backtracks up to the desired point but an alternative for the selected chapter is adopted and story continues from this point on. Notice that there is also a window listing the events corresponding to the selected chapter.

Besides the strategies for incorporating weak and strong interactions that we had in the first version of Logtell, further kinds of interactions are under consideration, such as letting users to: insert abstract events and situations in the story, which are automatically specialized by the simulation process; tune narrative tensions by means of numeric scales referring to levels of violence, romantic turns, etc.; share stories with other users, so that more than one user influence the same story; and suggest the insertion of events and situations using natural language.

5 Planning and Interaction at the Dramatization Level

The Drama Manager converts all events into actions, which are delegated to specific virtual actors, at specific times. In the first version of Logtell, for each event there was a single combination of actions for its representation. There was little diversity in the dramatization and the time for each event almost did not vary. In the second version, we introduced nondeterminism for the specification of each event. Each event is assigned to a nondeterministic automaton, created by authors and stored in the context database. By doing this, we try to increase the number of possible dramatizations, to enable the users to interfere in the story at this level, and to vary the time for the dramatization of an event, according to the convenience of the overall narration of the story. In each automaton, states are described by invariants, that is, logical formulae relating dramatization attribute variables. An automaton is said to be in a certain state if and only if its invariant is satisfied. Transitions correspond to actions delegated to actors. These actions are nondeterministic and can lead the automaton to alternative states. The new state may depend on the interference of the user or can be achieved according to the autonomy of the actors. Figure 3 shows an automaton created to represent possibilities for the dramatization of an event, in which a villain kidnaps a victim.

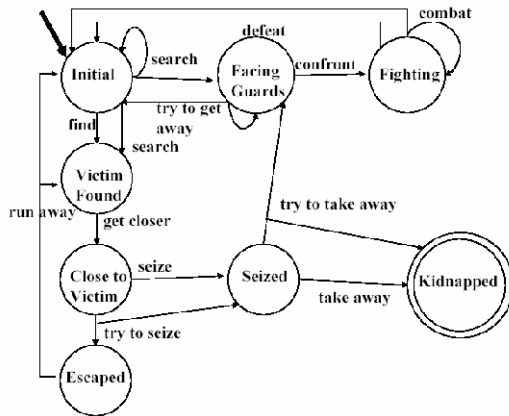


Fig. 3. Example of automaton for the representation of an event

The Drama Manager has now an abstract control layer that controls the automaton corresponding to the current event. The control layer checks the values of the attribute variables and detects the current state. Since the current state is known, an action is chosen to be performed by the actors. The choice of the action depends on alternative plans for leading the automaton from the initial state to a final state. These plans are called policies and they map states into actions to be performed.

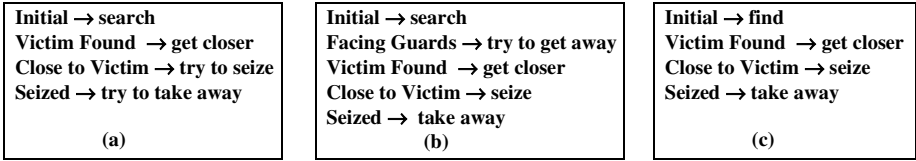


Fig. 4. Examples of policies: weak (a), strong with cycles (b) and strong (c)

In order to create policies, we adopt planning as model checking [19, 20], which is an approach for planning under uncertainty. Model checking [21] is a formal method to check whether a certain logical formula is a model for a structure. Reachability of a certain state in an automaton is a typical application of model checking. In planning as model checking, the planning algorithm examines the automaton in order to create policies that take the automaton from one state to another state. Generated policies can be weak, strong with cycles or strong. In weak policies, there is at least one path from the initial state to the goal state, but states from which it is not possible to reach the goal according to the policy can be reached. When a policy is strong with cycles, the goal state is always reachable but cycles can occur, so that the time to reach the goal might be virtually infinite. Strong policies guarantee that, from any state in the policy, the goal state is reached at some moment. A state is said to be safe if it has at least one strong policy leading the automaton from this state up to the final state. In our automata, only safe states are allowed. Figure 4 shows weak, strong with cycles and strong policies for our example.

Based on the automaton for an event, various policies of all kinds can be built beforehand using planning as model checking. In this way, there is no need to spend time planning during dramatization. The set of policies correspond to alternatives for leading the automaton from specific states to the final state. Occasionally, policies to postpone the end of an event, leading the automaton to intermediate states might be useful. Policies are chosen and replaced during dramatization in accordance with some principles: (i) In order to obtain variety during dramatization, the choice should occur at random up to a certain extent. (ii) At the beginning of the dramatization of an event, weak policies are accepted (and even preferred) because they tend to be more realistic, there are usually many weak policies available (which guarantees diversity) and duration is not yet a problem. When a state not mapped to any action in the current policy is reached, the policy has to be replaced. (iii) At a certain point of the dramatization, there might be time to accept cycles, but it might not be convenient to postpone termination for a long time. Unnecessary change of policies should be avoided and policies that are strong with cycles are preferred. (iv) When there is little time available, strong policies are needed to force termination.

The main difficulty for user interference in the dramatization is the fact that the scene might be modified in such a way that the story would become inconsistent. By using planning as model checking, it is possible to introduce opportunities for interaction at this level. If the nondeterministic automaton considers opportunities for user intervention and this intervention is restricted to automaton's transitions, there is no risk of not reaching the final state of the current event. Forms of interaction at the dramatization level are still under investigation, but two alternatives seem to be straightforward. The first one corresponds to interactions in which the user would choose the result of a nondeterministic action explicitly. In the second one, users would control actions of avatars, but restricted to possibilities specified by the automaton.

Besides the use of automated planning for plot composition and for controlling the dramatization, our virtual actors have a minimum of planning capabilities, at a low level of detail. Since actors are expected to play the assigned roles, and must plan in order to achieve an adequate performance, some simple planning resources become indispensable, so that, in real-time, an actor be able to make decisions and to schedule the necessary micro-actions. In general, simple path-finding algorithms and direct inter-agent communication schemes are sufficient. Each actor must also incorporate behaviours for interacting with the physical environment and with the other actors. The local planning of each actor must be simplified to ensure short response times.

6 Concluding Remarks

Logtell has been extended to work in an interactive TV environment, conciliating requirements of coherence and diversity of stories with high responsiveness. In this paper, we presented the various levels of automated planning and interaction we have incorporated to perform tasks that, in filmmaking, are typical of authors, directors and actors. Planning techniques have shown to be important to provide diversity and coherence for our stories. In order to cope with high responsiveness, new interaction methods have been investigated and implemented and, as efficiency became essential, the planner for plot generation has been enhanced. Planning as model checking has also been introduced in dramatization, which is important to diversify the ways events are dramatized, to control dramatization time and to let users interact with the story at the dramatization level.

Several other research topics are being investigated to achieve realism and good quality in the narration of stories by Logtell, including models for representing beliefs and emotions of characters, alternative ways to control the camera, and automatic text generation for supporting dialogs and the explicit narration of stories.

References

1. Cavazza, M., Charles, F., Mead, S.: Character-based interactive storytelling. *IEEE Intelligent Systems*, special issue on AI in Interactive Entertainment 17(4), 17–24 (2002)
2. Grasbon, D., Braun, N.: A morphological approach to interactive storytelling. In: *Proc. CAST 2001, Living in Mixed Realities*. Special issue of *Netzspannung.org/journal*, the Magazine for Media Production and Inter-media Research, Sankt Augustin, Germany, pp. 337–340 (2001)

3. Paiva, A., Machado, I., Prada, R.: Heroes, Villains, Magicians... Dramatis personae in a virtual story creation environment. In: Proc. Intelligent User Interfaces (2001)
4. Propp, V.: Morphology of the Folktale, Laurence Scott (trans.). University of Texas Press, Austin (1968)
5. Mateas, M., Stern, A.: Towards integrating plot and character for interactive drama. In: Dautenhahn, K. (ed.) Socially Intelligent Agents: The Human in the Loop, AAAI Fall Symposium, Technical report, pp. 113–118. AAAI Press, Menlo Park (2000)
6. Crawford, C.: Assumptions underlying the Erasmatron storytelling system. In: Working Notes of the 1999 AAAI Spring Symposium on Narrative Intelligence. AAAI Press, Menlo Park (1999)
7. Riedl, M., Young, R.M.: An intent-driven planner for multi-agent story generation. In: The Proceedings of the 3rd International Conference on Autonomous Agents and Multi Agent Systems (2004)
8. Ciarlini, A.E.M., Pozzer, C.T., Furtado, A.L., Feijo, B.: A Logic-Based Tool for Interactive Generation and Dramatization of Stories. In: Proc. ACM SIGCHI International Conference on Advances in Computer Entertainment Technology (ACE 2005), Valencia (2005)
9. Ciarlini, A., Veloso, P., Furtado, A.: A Formal Framework for Modelling at the Behavioural Level. In: Proc. The Tenth European-Japanese Conference on Information Modelling and Knowledge Bases, Saariselkä, Finland (2000)
10. Yang, Q., Tenenber, J., Woods, S.: On the Implementation and Evaluation of Abtweak. *Computational Intelligence Journal* 12(2), 295–318 (1996)
11. Erol, K., Hendler, J., Nau, D.S.: UMCP: A sound and complete procedure for hierarchical task-network planning. In: Proceedings of the International Conference on AI Planning Systems (AIPS), pp. 249–254 (1994)
12. Nau, D.S., Au, T.-C., Ilghami, O., Kuter, U., Murdock, W., Wu, D., Yaman, F.: SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research* 20, 379–404 (2003)
13. Blum, A.L., Furst, M.L.: Fast Planning through planning graph analysis. *Artificial Intelligence* 90, 281–300 (1997)
14. Geverini, A., Serina, I.: LPG: A planner based on local search for planning graphs. In: Proceedings of the International Conference on AI Planning Systems (AIPS), pp. 968–973 (2002)
15. Doherty, P., Kvarnström, J.: TALplanner: A temporal logic based planner. *AI Magazine* 22(3), 95–102 (2001)
16. Hoffman, J.F.F.: The Fast-Forward planning system. *AI Magazine* 22(3), 57–62 (2001)
17. Aarne, A.: The Types of the Folktale: A Classification and Bibliography. In: Translated and enlarged by Stith Thompson, FF Communications, vol. 184. Suomalainen Tiedekatemia, Helsinki (1964)
18. Kautz, H.A.: A Formal Theory of Plan Recognition and its Implementation. In: Allen, J.F., et al. (eds.) Reasoning about Plans. Morgan Kaufmann, San Mateo (1991)
19. Bertoli, P., Cimatti, A., Roveri, M., Traverso, P.: Planning in Nondeterministic Domains under Partial Observability via Symbolic Model Checking. In: 17th International Joint Conferences on Artificial Intelligence, Washington, pp. 473–478 (2001)
20. Pistore, M., Traverso, P.: Planning as model checking for extended goals in nondeterministic domains. In: Proc. 17th International Joint Conference on Artificial Intelligence, Washington, pp. 479–484 (2001)
21. Burch, J.R., Clarke, E.M., McMillan, K.L., Dill, D.L., Hwang, L.J.: Symbolic Model Checking: 10^{20} states and beyond. *Information and Computation* 98(2), 142–170 (1992)

Exploiting Structure and Conventions of Movie Scripts for Information Retrieval and Text Mining

Arnav Jhala^{1,2}

¹ Center for Computer Games Research, IT University of Copenhagen,
7 Rued Langgaards Vej, Copenhagen, Denmark

² Digital Games Research Center, North Carolina State University
890 Oval Dr, Raleigh, NC - 27606
ahjhala@ncsu.edu, arjh@itu.dk

Abstract. Movie scripts are documents that describe the story, stage direction for actors and camera, and dialogue. Script writers, directors, and cinematographers have standardized the format and language that is used in script writing. Scripts contain a wealth of information about narrative patterns, character direction, blocking, and camera control that can be extracted for various applications in interactive storytelling. In this short paper, we propose the creation of an automatically annotated corpus of movie scripts and describe our initial efforts in automating script annotation. We first describe the parts of a movie script that can be automatically annotated and then describe the use of an existing language processing toolkit to automatically annotate specific parts of a movie script.

Keywords: Text mining, Automatic Annotation, Narrative Patterns.

1 Introduction

Automated story generation has been an active area of research in recent years. Advances in and popularity of computer games technologies have opened up new forms of storytelling for entertainment and educational applications. Most automated storytelling systems are based on cognitive models of storytelling that are implemented through symbolic approaches. Movie scripts can provide a wealth of information about story structure and relationships between actions in the story and camera directives. Movie scripts follow a standard set of terms and formatting conventions [1]. These formatting conventions and standard terms can be used to automatically extract information about different elements of scripts. In this paper I describe an ontology of terms commonly used in movie scripts. This ontology is then used in conjunction with annotation rules over a corpus of scripts for automatic annotation using the GATE natural language toolkit [2]. The advantage of annotating scripts automatically is that annotated scripts can be easily queried and analyzed. For instance, the number of occurrences of wide angle camera shots in a particular genre of movies can be obtained as a parameter for classifying movies automatically based on script. Wide-angle shots are used with more frequency in adventure and war movies than in romantic movies, Wide angle shots are specified either as 'WIDE' or

‘WIDEANGLE’ camera directive and are usually capitalized. The capitalization and position of occurrence is used for annotating it as a camera directive.

2 Anatomy of a Movie Script

The script is a communicative tool that is used throughout the movie development pipeline. The movie is an audiovisual communicative medium and the script represents audio-visual communicative intent of the writer and the director. Screenwriting resource, screenwriting.info aptly defines the script and elaborates the script writing process as: “A script is a document that outlines every aural, visual, behavioral, and lingual element required to tell a story. Because film is a highly collaborative medium and the director, cast, editor, and production crew will, based on an outline, interpret the story their way when it is filmed. But because so many people are involved in the making of a film, a script must conform to standards that all involved parties understand and thus has a specific format or layout, margins, notation, and other conventions.” The script, on the basis of content and formatting, can be divided into the following sections:

- Title
- Author name(s)
- Scene Heading
- Shot
- Transition
- Action
- Character Name
- Dialogue
- Parenthetical
- Extensions

The title of the movie is at the top of the script followed by the author. Scene heading establishes the physical context of the action that follows it. Scene heading usually includes an EXT or INT directive for external or internal location followed by the name of the location and time of day. Sometimes scene heading also includes comments like ESTABLISHING SHOT. An example of a scene heading is: “EXT. KEY WEST MARINA - DAWN – ESTABLISHING”. Camera directions like FADE IN/OUT are also included in the scene heading. The scene heading, which establishes the physical and temporal context, is followed by the description of action taking place at the location. In case of an establishing shot, important elements of the location are described through panning shots or individual shots. Ambient sounds are also described in the establishing shots of the location. Actions of characters are described along with the dialogues for each character. Dialogues have the character name (usually in all caps) followed by the dialogue text. In most scripts, dialogue actions are clearly separated from action and camera directions by enclosing the non-dialogue descriptions in brackets. Parenthetical directives are not commonly used but

they are meant to describe phenomena such as inflection or mood. For instance, sleepily in the following dialogue segment is a parenthetical.

JULIE (*sleepily*)
What? What time is it?

O.S. – Off Screen and V.O. – Voice Over are examples of extensions to the dialogue frame that indicate if the voice is meant to come from off screen or if it is a voice over of an extra-diegetic narrator. Transitions are described using standard terms and are usually capitalized. CUT TO, PAN TO, FADE TO, DISSOLVE TO, and FADE OUT are common transitions. Finally, shots indicate the type of shot and point of view of the camera. Examples of shot types are CLOSE UP, POINT OF VIEW, REVERSE SHOT, etc.

3 Annotating Movie Scripts

The General Architecture for Text Engineering (GATE) [2, 3] is a software developed at the University of Sheffield for various natural language processing, information retrieval and text mining applications. GATE is a plugin-based architecture developed in Java [3]. There are three different types of components that are used as resources by GATE. Components are software modules with well-defined interfaces for interaction. The three types of resources that are used in GATE are:

1. Language Resources (LRs) represent documents, ontologies and corpora that are input to the system.
2. Processing Resources (PRs) represent applications that are part of the NLprocessing pipeline like parsers, POS taggers etc.
3. Visual Resources (VRs) are user interface elements that are used for editing or browsing the LR and PRs. Ontology editor and document editor are examples of Visual Resource plugins.

GATE comes with pre-written processing resources included in a package called ANNIE (A Nearly New Information Extraction System). In order to make the IR task on movie scripts more efficient and informed the standard terminology and structure of movie scripts can be utilized. The GATE toolkit is used for creating an automatically annotated corpus of movie scripts with annotations specifying the different elements of the scripts described in the earlier section. Scripts are annotated by using ANNIE's processing resources to identify Names, Dialogues, Camera Direction, Actions, etc. The Onto-Gazetteer in ANNIE is given as an input an ontology file containing standard terms for camera and stage directions as described in the earlier section. On detecting the standard term, ANNIE annotates the term using rules from an input file to one of the sections of the document. The inputs to the component are in the form of .lst files that list out terms, for instance, camera.lst file contains standard terms for the camera like Establishing Shot, Wide-Angle, etc. The rules file contains rule such as camera.lst:sc:direction which associates the terms in the camera.lst file with the direction annotation.

4 Conclusion

Information extraction on movie scripts can yield a wealth of information about dialogue, cinematic conventions, narrative structure, etc. In this paper, I described an automatic annotation tool that takes into account the structure and conventions of movie scripts for language processing applications. The tool is developed on an existing language processing toolkit and is integrated with the other components of the NL pipeline. I developed an ontology of semantic parts of the script based on filmmaking and script writing standards. The ontology overlaps with semantic multimedia annotation standards [4] being developed by the semantic web community and is hence potentially a contribution to the semantic video annotation standards. A corpus of annotated scripts can be a valuable resource for retrieving information about camera shots and story structure and applying statistical techniques on the annotated corpus for various other applications. While much work still needs to be done in creating a corpus of annotated scripts for analysis, this initial effort is a successful step towards having a large scale corpus of annotated scripts.

References

1. Screenwriting Tutorial, <http://www.screenwriting.info>
2. Cunningham, H.: GATE-A General Architecture for Text Engineering. *Computers and the Humanities* 36, 223–254 (2002)
3. Cunningham, H., Humphreys, K., Gaizauskas, R., Stower, M.: CREOLE Developer's manual, Technical Report, Department of Computer Science, University of Sheffield (1996)
4. Jewell, M., Lawrence, F., Tuffield, M., Bennett, A., Millard, D., Nixon, M., Schraefel, M., Shadbolt, N.: Ontomedia: An Ontology for the Representation of Heterogeneous Media. In: *Multimedia Information Retrieval Workshop*. ACM SIGIR, Brazil (2005)

Generation of Dilemma-Based Narratives: Method and Turing Test Evaluation

Heather Barber and Daniel Kudenko

University of York, Heslington, York, YO10 5DD, England
{hbarber,kudenko}@cs.york.ac.uk

Abstract. This paper describes the analysis of the narrative generation aspect of the Generator of Adaptive Dilemma-based Interactive Narratives (GADIN). This system automatically generates narratives that are focused around dilemmas to create dramatic tension. The dilemmas are based on those clichés encountered in many of today’s soap operas. A story planner creates sequences of actions that lead to dilemmas for storyworld characters. A narrative generated by GADIN is compared to that which took place in a television soap opera. The stories were subjected to a Turing-style test, with 127 participants, which showed that the readers could not determine which narrative had been written by a computer.

1 Introduction

In this paper we analyse a system that automatically generates narratives. This is one aspect of the Generator of Adaptive Dilemma-based Interactive Narratives (GADIN). Dramatic tension is created by incorporating dilemmas for storyworld characters. These dilemmas are based on the clichés found in many contemporary soap operas, such as the trade-off between personal gain and loyalty to a friend. Overarching stories connect these dilemmas within a coherent plotline that is dynamically created.

Story generation systems [1,2,3,4] created thus far have written stories which have clearly not had a human author. The stories generated by the GADIN system are more consistent and have a level of dramatic interest equivalent to that in human-authored stories. This is demonstrated here through a Turing-style test, in which participants were unable to distinguish the GADIN narrative from one which had occurred in a television soap.

Soap operas (or soaps) are a popular means of entertainment. They consist a story which never ends, shown in episodes which see the continuation, foundation or conclusion of some of the ongoing storylines. Soaps are commonly known to consist of very similar and clichéd storylines. The clichéd storylines utilised in soaps are usually found to be essentially conflicts (or dilemmas). Writers utilise these dilemmas in the creation of soap stories in which the story is built around the cliché. A general form of each such clichéd dilemma can be determined, and GADIN can generate a soap narrative around these. The soaps which are particularly concentrated on in the discussed version of GADIN are *Hollyoaks*, *Eastenders*, *Coronation Street*, *Emmerdale* and *The Archers*.

2 GADIN

Many storytelling genres make frequent use of clichéd storylines which are created around dilemmas to storyworld characters. These dilemmas can be generalised and the GADIN architecture uses planning to achieve such dilemmas, the combination of plan and dilemma constitute a dramatically interesting sub-story of the generated narrative. Characters act and experience these dilemmas in the course of a GADIN narrative, making decisions and action choices depending on their individual properties and state.

More details on narrative generation in GADIN for soap operas can be found in [5]. The application of GADIN to children’s dinosaur adventures is discussed in [6].

3 Story Quality Evaluation

The Turing Test [7] can be interpreted as requiring a human reader to be unable to determine whether a narrative was written by a computer or human author. In the test described here the reader must decide which of two narratives they think was written by a computer.

To ensure a fair comparison both narrative creations began with the same storyworld state. The characters were limited to only those included in the selected story, but the actions and dilemmas included all of those available to the GADIN system. Subjects who were familiar with the television soap narrative selected were asked to not answer the survey. The names were changed to be anonymous and the same in both narratives.

To make the narratives more readable they were adapted slightly. For example, the actions of a character moving between locations were removed in both versions – this is relevant in the experience but not in the subsequent telling. For the television soap it was necessary to transcribe the events in the form output by GADIN. Using this style of writing is an obvious limitation of the narrative quality, but in this evaluation only the core components of the narrative and its structure are required to be compared and thus this is sufficient. In both versions it was essential not to impose any reasoning on the characters in the wording.

The human-authored narrative used in this experiment has been taken from a television soap opera. The GADIN-authored narrative was not specially selected but was the first to be generated from the given storyworld state. It was ended when at the same length as the television soap narrative, as although the narrative generation would continue indefinitely from this point the later narrative content was not relevant to this comparison.

4 The Narratives

The narratives shown to the participants are given here.

STORY 1: Jane and Tom are in a relationship. Jane becomes pregnant. She decides not to keep the baby. Tom proposes but Jane rejects him, ending the

relationship. Nick and Sally go to the shop. Nick flirts with Tom. Sally flirts with Tom. Tom must choose whether to partner Sally or Nick. He decides to go out with Sally. Nick no longer fancies Tom. Nick starts to go out with Rich. Sally wants to start an affair with Rich but he chooses not to. Tom expresses his disapproval of Rich's relationship with Nick, but Rich decides to continue the relationship anyway. Jane and Sally stop liking one another. Tom starts an affair with Jane.

*STORY 2: Tom and Jane are in a relationship. Rich and Sally flirt with each other, and start going out. Rich wants to start an affair with Tom but he chooses not to. Rich splits up with Sally. Tom stops liking Rich. Rich and Sally start going out with each other again. Rich ends the relationship with Sally. Rich starts a relationship with Nick. Tom expresses his disapproval of this relationship but Rich ignores him. Rich and Tom start having an affair. Tom proposes to Jane and she accepts. Rich tells Jane about his affair with Tom. Jane ends her relationship with Tom. Tom and Rich start a relationship*¹

5 The Results

The survey was divided into two groups of participants. One consisted of those who regularly view soaps (although not the soap in question). These participants were targeted through posts on 4 English soap forums.² Given the focus on English soaps throughout it was important that these forums were used only by an English audience, as the style and content of soaps does vary to some extent between nationalities. The second group contained those who regularly play computer games and are thus more familiar with computer-based stories. For this the survey was posted on 2 English games forums and 1 international.³

It is generally considered that a result which has a more than 5% chance of occurring if random choices are made has successfully passed the Turing test. The results of this test were as follows:

- 42 soap viewers participated in the survey. Of these, 24 – or 57.1% – thought that story 1 was the one which had been written by a computer. This will occur 22% of the time with random choices.
- 85 games players took part in the survey. In this case 49 participants, or 57.6%, thought that story 1 was written by a computer. This will occur 10% of the time if random choices are made.
- Overall there were 127 surveys completed. 57.5% thought that story 1 had been written by a computer. This will occur more than 5% of the time if random choices are made.

¹ Story 1 was written by the GADIN system and Story 2 is from Hollyoaks.

² The soap forums utilised were at: www.soapforum.co.uk; www.digitalspy.co.uk; www.soapboards.co.uk; www.tvforum.co.uk

³ The games forums used were: www.enemydown.co.uk; www.na-clan.co.uk; www.wcreplays.com

From these results it can be concluded that the GADIN system is capable of generating a narrative which is not discernibly different from an outline of a story occurring in a television soap opera.

6 Conclusions

We have shown that this story generation method generates narratives which are indistinguishable from those in standard soap operas.

In this experiment only a single storyline has been evaluated. The infinite nature of soaps means that there will be a number of such storylines. However since this was chosen randomly the implication is that other storylines are likely to also be indistinguishable from those in television soaps.

This experiment is focused on the very generalised story development aspect and on the creation of the core component of dramatic interest and the underlying structure of the narrative. This will inevitably affect users opinions of the system and their experience but hopefully the fundamentals of the narrative maintain sufficient dramatic interest for this to not too adversely affect the results.

It would potentially be possible to use GADIN in collaboration with script writers to create televised soaps. Future experiments could involve the GADIN story outline being written as a script which would be performed by human actors. This could then be compared to the television soap script being acted by the same characters. This would compare the whole viewing experience as opposed to the core components of the narrative as was the focus in this experiment.

Acknowledgements

Thanks are due to everyone who participated in the survey, and in particular to Christian Wood who coordinated the games survey.

References

1. Meehan, J.R.: *The Metanovel: Writing Stories by Computer*. Garland Publishing, Inc., New York (1980)
2. Bringsjord, S., Ferrucci, D.A.: *Artificial Intelligence and Literary Creativity: Inside the Mind of BRUTUS, a Storytelling Machine*. Lawrence Erlbaum, Mahwah (2000)
3. Lebowitz, M.: *Planning stories*. In: *Ninth Annual Conference of the Cognitive Science Society*, Seattle WA (1987)
4. Turner, S.R.: *The Creative Process: A computer model of storytelling and creativity*. Lawrence Erlbaum Associates, Mahwah (1994)
5. Barber, H., Kudenko, D.: *Dynamic generation of dilemma-based interactive narratives*. In: *AIIDE 2007*, Stanford, California (2007)
6. Barber, H., Kudenko, D.: *Generation of dilemma-based interactive narratives with a changeable story goal*. In: *2nd International Conference on Intelligent Technologies for Interactive Entertainment*, Cancun, Mexico (2008)
7. Turing, A.: *Computing machinery and intelligence*. *Mind* 59 (1950)

Emergent Stories Facilitated

An Architecture to Generate Stories Using Intelligent Synthetic Characters

Rui Figueiredo¹, Antonio Brisson¹, Ruth Aylett², and Ana Paiva¹

¹ INESC-ID, Avenida Prof. Cavaco Silva - Taguspark, 2780-990 Porto Salvo, Portugal

² Heriot-Watt University, Edinburgh EH14 4AS, UK

rui.figueiredo@tagus.ist.utl.pt, antonio.brisson@gaips.inesc-id.pt,
ruth@macs.hw.ac.uk, ana.paiva@inesc-id.pt

Abstract. In this paper we address the issue of how to transform a set of actions and behaviors generated by autonomous characters into a meaningful and interesting story. To do that, we consider the issues involved in developing an architecture that is generic enough to be applied to entertainment/educational applications that convey stories through the interactions between intelligent graphical characters and a user. As such, issues in organizing the emergent content are discussed and a generic story facilitator (a kind of game master) is described. Finally two story-telling applications that use the architecture are described demonstrating the generality of the approach.

Keywords: Autonomous agents, synthetic characters, interactive story-telling, narrative generation, interactive virtual environment.

1 Introduction

Educational games are increasingly popular as a way to teach people about a certain subject, reinforce development, or assist in learning a new skill. In particular, educational video games, sometimes referred to as a edutainment or serious games (because they mix education and entertainment), can be seen as an effective tool for creating rich contexts, reflective of the real world, that help to promote learning [2]. Traditional video games tend to limit user interactivity in order to guarantee the sequential unfolding of a predefined story. There are, however, several approaches to creating game-like systems that try overcome this limitation. Some have adopted a branching narrative structure, in which a finite number of pre-scripted paths result from a choice made at a specific decision point [5] [8]. In other cases, work has tried to cover the whole space of possible options [9], with a correspondingly combinatorial authoring problem.

All these approaches try to tackle the problem that occurs when users have the freedom to do what they please. The dilemma is that too much user freedom may result in interactions with the system that fail to produce a meaningful story-experience. However, constraining user interactivity too tightly produces a system that tells only one story, always the same, in which the user quickly

loses interest. This is especially true in an educational context where the stories are typically not so appealing to users as in traditional video games.

The type of application under investigation is a computer-based version of educational role-play. In educational role-play, social interaction is used as the stimulus for challenging and changing existing beliefs [14] and can result in significant behavioral changes [7] making it highly relevant for social and emotional learning [3] [6]. In role-play the story is improvised rather than scripted and emerges from interaction between the characters involved. Since the story is emergent there is no absolute guarantee that it will convey a meaningful message. This problem is handled in traditional educational role-play through a person, usually referred to as the facilitator of the role-play. This facilitator is responsible for ensuring that the story develops within certain boundaries relating to the objectives originally conceived for the role-play.

Role-play is often developed as a succession of episodes, in which the occurrence of external events as well as the consequences of actions within an episode may be controlled by the facilitator of the role-play. Each participant has a back-story that describes the role of his/her character, communicated by the facilitator at the start of each episode. In some cases the facilitators will themselves play a character with the specific intention of shaping the emerging story in particular ways. It is through the use of these methods that it is possible to bound the scope of the emerging story around a particular educational topic.

This paper proposes an architecture for a virtual version of a human facilitator that mimics the methods used by a real facilitator. In the remainder of this paper the proposed architecture is described and two examples of its application are given.

2 Related Work

Research in Interactive Storytelling applications not only raises technical and artistic challenges, but also specific challenges that derive from the need for cooperation between these fields. Indeed, the conflict between the user's interactional freedom and authorial intentions with respect to story development strongly relate to the artistic goal of structuring the story on the one hand and the technical goal of achieving a good user experience on the other. This has also been called the *boundary problem* [13].

The architectures that have been proposed to tackle this challenge can be divided into the *Character Centered Approach*, where all the knowledge of the story structure is in the character's AI, and the *Plot-based Approach*, in which a global action sequence forms the desired story. Both may implement an entity that mediates the possible conflicts between user actions and desired story development, though with very different intentions and outcomes.

An unmediated *Character Centered Approach* has been used in the Interactive Storytelling prototypes by Cavazza et al. [22] based in a very strong and detailed character AI. This approach achieved good evaluation metrics such as story scalability [23], nevertheless the characters AI is confined to the author's

definitions that end up guiding the user through a story limiting its interference in the outcome.

Projects that use *Mediation*, can implement several independent strategies to include the user in the story. An example is the MOE architecture proposed by Weyhrauch [17], in which the mediator distinguishes between acting upon the characters internal state, world state or the state of the interaction. In Mimesis [19] [20], acting within a plot-centred approach, *Mediation* is used to decide whether to *Accommodate* user/author conflicts by adjusting the author plans, or to *Intervene* causing a user conflicting action to fail. In Façade [11] [10] the mediation is more subtle because the story is fragmented into different *Action Beats* representing all possible outcomes, that can be interrupted and sequenced to keep up the story flow.

Using Mediation is a step toward including some user actions in the story, increasing immersion and agency. Nevertheless, in the above examples character autonomy is still very reduced limiting their reactions to environment changes, raising the possibility of making characters act ‘out of character.’ It also limits user interaction, and interference in the story outcome. Our approach the *Story Facilitator* (SF) described in the following section addresses this issue by implementing a *Mediation* architecture that is prepared to shape the outcome of an unscripted narrative by interacting with an environment populated by autonomous and affective characters.

3 Architecture

In educational role-play all the participants have a role which is communicated to them by the facilitator. Usually, that role is presented through a back-story (a description of the character from which the participant can interpret his/her role). After the participants have internalised their roles, the role-play scenes begin. It is part of the role of the facilitator to decide when each episode ends, and when the next one begins. During an episode, the facilitator might intervene if that seems necessary (for example, if the role-play is taking a path that is in conflict with that intended). The facilitator might even take the role of a character in order to guarantee that certain events happen if these are important for the intended message to be successfully conveyed by the role-play.

To achieve this behavior, the architecture that has been produced (Figure 1) contains a *Story Facilitator* agent, a *Story Memory* and a set of *Narrative Actions* that are available to the *Story Facilitator* (SF). The story is divided into episodes that roughly represent role-play scenes. Each episode is a description of where the action takes place, who is going to participate in it, what is the role of each participant, when it can be selected and when it should end, and a set of rules that describe in what situations the SF should intervene, and which action it should take during the intervention.

This architecture assumes that the environment is populated by autonomous agents and user(s) whose actions produce events that can be perceived by the SF agent. The SF agent can itself change the environment through a set of actions named *Narrative Actions*.

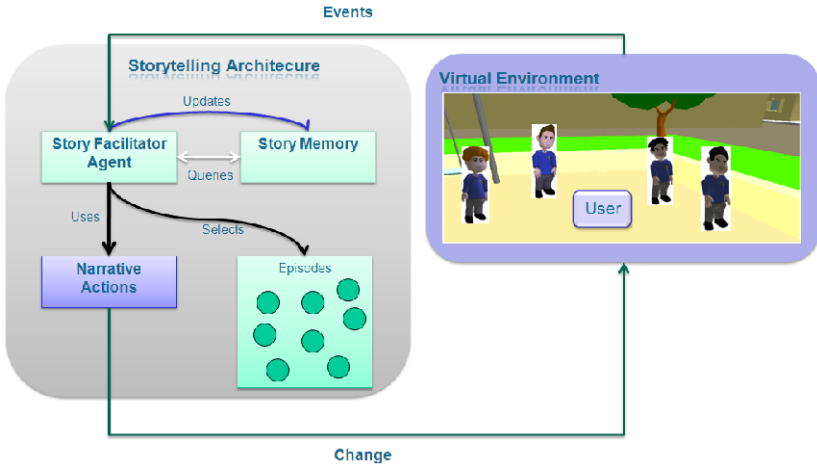


Fig. 1. The architecture

3.1 Episodes

An episode (Figure 2) represents a part of the story that can be combined with other episodes, where each combination produces a different overall story.




Fig. 2. Episode Structure

As the story unfolds and the SF perceives the actions that occur in the environment it populates the *Story Memory* with *Story Events* - descriptions of the actions the characters performed. The episodes' preconditions are matched against the *Story Events* and when satisfied indicate that the episode is a candidate for selection by the SF. The *Story Memory* also contains information that describes the characters, for example whether a character is hurt or not. This information is stored as *properties* and can also be matched against the preconditions of episodes.

The finish conditions determine when the SF should end an episode and select another for the continuing story and are identical in form to the episode preconditions. Finish conditions are required because in a character-based approach

there is no pre-determined sequence of actions to work through from start to finish and more abstract criteria such as the satisfaction of specific goals may be required for an episode to be considered finished.

Each episode contains a list of the participants that take part in it. A participant can be a character or an object(prop). For each character participant the episode contains a list of its goals and properties. This can be seen as a explicit representation of their back-story, communicated to them when the episode starts. A property for a participant is a characteristic inherent to him/her/it. For example, a property relating to an object (prop) representing a book might be called *belongs-to* and have as its value the name of one of the characters. For character participants the episode contains a list of goals for each of them; these are held as goal names, with the actual implementation of the goal located in the autonomous agent that controls the character.

Each episode has an intro/setup which is a set of *Narrative Actions* the SF should perform as soon as the episode starts. The typical intro/setup for an episode is the placing of the characters in the environment.

Similarly to what happens in real educational role-play there might be situations where the SF has to intervene during an episode. Those situations are specified as *Triggers*. A trigger is composed of a condition and a set of *Narrative Actions* that the SF will perform if the trigger's condition is satisfied.

3.2 Story Memory

The *Story Memory* represents the story as it is perceived by the SF agent. Every time a character performs an action (be it a character controlled by an autonomous agent or a user) the SF perceives the action and creates a *Story Event* (Table 1) that it stores in the *Story Memory*.

Table 1. Attributes of a story event

Attribute	Description
Subject	The ID of the agent (or user) that performed the action
Action	The name of the action
Parameters	The parameters used in the invocation of the action
Event Type	Indicates if the action is starting or finishing
Episode	The episode where the story event took place

Also, each time a character is placed in the virtual environment the SF records in the *Story Memory* the character's properties (Table 2) contained in the selected episode's definition.

Conditions. Using the information stored in the *Story Memory* the SF tests if the conditions in the preconditions, finish conditions or triggers contained in the episodes are satisfied. The episode's preconditions, finish conditions and the triggers' condition consist of conjunctions and/or disjunctions of *Event* and

Table 2. Attributes of a property

Attribute	Description
Name	The name of the property
Value	The value of the property
Holder	The character or object to whom this property refers to

Table 3. Query to a Story Event

Field	Description
Subject	The id of a participant(character, object or user)
Action	The name of the action
Parameters	The parameters for the action being tested
Episode Name	The name of the episode where the event was generated
Event-Type	If the event is an event that indicates that the action is starting or is finishing
Negated	If true indicates that the condition is true if there is no story event that satisfies it

Table 4. Query to a Property

Field	Description
Property Name	The name of the property
Holder	The name of the holder of the property, can be an object, an agent or the user
Value	The value of the property
Operator	One of the following operators: <i>Equal, NotEqual, LesserThan, GreaterThan</i>

Property Conditions, which are the two types of conditions that the SF can test.

An *Event Condition*(Table 3) tests if there is any (or no) *Story Event* in the *Story Memory* that satisfies it. All the parameters of an *Event Condition* are optional. If the fields: *subject*, *action*, *parameters* and *episode name* are not specified this indicates that they may take any value. If not specified, the *event-type* parameter indicates that we are testing for an event that indicates that the action has finished. For instance, consider a condition that only specifies the *subject*, for example *John*, and the action *Cry*. This condition will always return true after the character *John* finishes the action *Cry*.

The other type of condition is the *Property Condition*(Table 4). In this type of condition all the fields except the field *Operator* can be omitted, indicating that their value should not be considered. If the field, *Value*, is omitted the *Operator* field is ignored. Imagine as an example, a query to a property where we specify the *property name* and the *holder* and we omit the *value*. This query returns false only when there is no such property for that *holder*.

3.3 Narrative Actions

The *Narrative Actions* (Table 5) are the actions the SF has available to intervene while the story is unfolding. They are used during the *intro/set-up* of the episodes and when a *trigger* is selected for execution.

Table 5. Narrative actions available to the author

Narrative Action	Description
Load Scenario	This action loads a scenario (e.g. 3D model of a classroom)
Insert Character	This action inserts an agent in the current episode
Insert Object	This action inserts an object in the current episode
Narrate	Writes text on the interface. Used to simulate a narrator
Add goal	Adds a goal to a particular character
Remove goal	Removes a goal from a particular character
Act for Character	Makes a specific character perform an action
Remove Object	Removes an item from the set
Remove Character	Removes a character from the set

The *Insert Character/Object* are mainly used during the *intro/setup* of an episode to place the characters and objects in the environment. They can also be used while an episode is unfolding through the use of a trigger, for example to insert a new character when a specific set of events occurs.

The *Add/Remove goal* actions are used to change the behavior of a character when certain events occur. These *Narrative Actions* can be used as “memory packets” as described in [26]. In role-play a *memory packet* is something a character is not to remember until an appropriate time in the game. Using an example taken from [26] that refers to a role-play about *Hamlet*: “...a (memory) packet meant to be opened after drinking a special kind of tea, a packet that might tell her (the character) that ... she is now passionate about botany and has forgotten all about Hamlet...”. The *Add/Remove goal* action allow us to write triggers for the episodes that behave as “memory packets” as they are used in real role-play.

The *Act For Character* action is used to allow the SF to take control of a character. It could be used to ensure that one or more events take place in the story so that the message to be conveyed by the role-play is understood. Furthermore, because this architecture is meant to be used with autonomous agents, which require a great deal of authoring, the *Act For Character* action may be used to create scripted behaviors, for example, to make a character walk to a particular part of the set and introduce himself, that otherwise would be difficult to create with goals and plans in an agent architecture based on planning.

3.4 Triggers

As the story unfolds there may be situations where the actions of the characters and/or user(s) alone are not enough to progress the story. An example would be

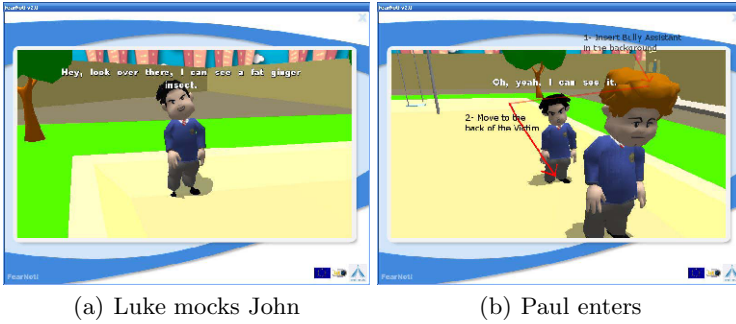


Fig. 3. The trigger is fired when Luke mocks John

where another character is supposed to enter after a particular interaction between the characters already in a scene. To achieve such exogenous behaviors, we (as authors) specify the triggers that signal the moments when SF’s intervention is necessary during the story when creating the episodes. For each trigger it is necessary to specify a condition (that states when the *trigger* should be “fired”) and a set of *Narrative Actions* to be executed when the trigger is selected for execution.

As an example of a *trigger*, imagine a situation where we want a character to make an appearance (*Paul*) when another (*Luke*) mocks the victim (*John*) (Figure 3). That *trigger* has as its condition that the character *Luke* performs the action *Mock* to the character *John*, and as the *trigger’s* narrative actions it has the *Insert Character* narrative actions, followed by an *Act For Character* narrative action that will make the inserted character move to a particular point of the set (behind *John*).

3.5 Story Facilitator (Cycle)

The Story Facilitator’s tasks can be described in a cycle (Figure 4) that is repeated until the story ends. At the very beginning (*Start*) the SF loads all the episodes for the story and puts them in the *Story Memory*. It then selects an episode that has an empty set of *preconditions* and marks that episode as *selected* so it will not be selected again. The SF will then execute all the *Narrative Actions* contained in the selected episode’s *intro/set-up*. Each time a character that is not controlled by a user is inserted on the set, the SF reads that character’s *properties* and stores them in the *Story Memory*, and then sends the character’s goals listed in the selected *episode* to the agent that is responsible for controlling the character. After the *intro/set-up* of the episode ends, the SF hands control to the agents and user(s) controlling the characters, making the story emergent from that point on.

During this phase, where the agents and user(s) control the characters present in the set, the SF perceives all actions performed in the virtual environment and updates the *Story Memory* with *Story Events*. If the content of the *Story*

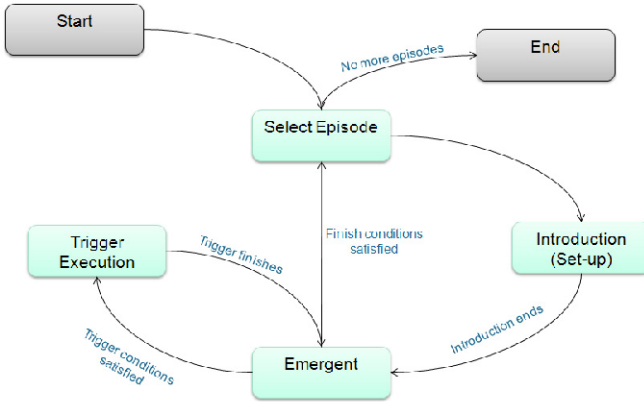


Fig. 4. The stages the episodes go through

Memory satisfies a *trigger's* condition, the SF intervenes and executes the *trigger's Narrative Actions*. Finally, if one of the episode's *finish conditions* is satisfied by the contents of the *Story Memory* the SF ends the current episode and uses the *Story Memory* to check if there are other episodes whose *preconditions* are satisfied. If so, the SF selects one of them as the next episode. If there is no other episode that can be selected, the SF ends the story.

4 Case Studies

4.1 FearNot!

FearNot! (Figure 3) is an Interactive Virtual Environment (IVE) developed to be used as an educational tool to promote awareness about bullying behavior in schools. *FearNot!* is inspired by role-play. Each character has a role that is related to a typical bullying scenario in a school. There are bully characters, bully assistants, victims, defenders (who try to help the victim), bystanders, etc. The user has the role of friend of the victim. The story develops as the user is asked for advices by the victim character on how to cope with the bullying situations he or she is experiencing. One version has only male characters and the other only females. The victim tries to follow the advice, although sometimes they might not be able to.

The agents that control the characters utilise the FATiMA agent architecture [15] [16] based on the OCC model of emotions [12]. Their goals depend on their emotional state, making the user aware that the actions characters carry out and whether those actions succeed or not is related to how they feel, not just to logic. This is particularly important in the case of the victim character, who often “feels” too much fear to comply with some of the user’s suggestions, for example to fight back against the bully.

FearNot! has roughly 40 episodes that can be selected by the SF and has just been evaluated in English and German schools.

4.2 I-Shadows

I-Shadows is an Interactive Storytelling application that takes the form of a Chinese Shadow Theater. In this theater the user interacts with the system by controlling a physical puppet that is either a Hero or a Villain and whose movements are interpreted by a vision system that sends that information to the autonomous character controllers. This system is inspired by fairy tales and the characters are authored approximately according to the morphology of Vladimir Propp. The story develops as the user interacts with the other characters using expressive movement [25]. Because I-Shadows also implements the same autonomous agents and Story Facilitator as in the FearNot! case, it allows the system to measure the emotional development of the scene and as a consequence the moments of tension and resolution [24] of the story. Using this information and the relations between characters, the system can carry out actions in order to shape the story flow by adding or removing characters and objects that contribute to a predefined mood development. I-Shadows is currently in its user testing phase, however the experience of implementing the Story Facilitator in the system as been very encouraging and successful.

5 Conclusions and Future Work

In this paper we described an architecture inspired by educational role-play that can be applied to story-generating environments populated by intelligent synthetic characters. We describe two applications where this architecture was used, *FearNot!* and *I-Shadows*. We believe that this architecture can easily be applied to applications that use intelligent synthetic characters in a virtual environment to create stories, as are the two systems we described.

The current version of the architecture requires a substantial amount of authoring work that could be reduced if planning abilities were added to the Story Facilitator (SF). In traditional role-playing games, such as *Dungeons&Dragons*, an outline story is planned by the *game-master* prior to the actual role-play by the characters [1]. As the role-play develops the *game-master* may intervene in order to shape the story according to the initial conception, for example by changing the environment, by determining the outcome of user actions, by taking control of a character, etc. If the SF agent is capable of processing story goals and planning story unfolding, the triggers the author has to write to account for unwanted character actions that might damage the story, could be handled automatically in a mediation process similar to the one used in [18] [21]. Furthermore, if the episodes were annotated with episode-level goals, the SF could automatically decide when to finish the episodes, eliminating the need for specific finish conditions.

Acknowledgments. This work was partially supported by the European Community (EC) under the IST RTD programme, contract IST-2001-33310-VICTEC, and is currently funded by the eCircus project IST-4-027656-STP. The authors are solely responsible for the content of this publication. It does not represent the

opinion of the EC, and the EC is not responsible for any use that might be made of data appearing therein.

References

1. Aylett, R.S., Louchart, S., Tychsen, A., Hitchens, M., Figueiredo, R., Delgado Mata, C.: Managing Emergent Character-Based Narrative. In: Proceedings, IN-TETAIN 2008, Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering (January 2008)
2. Bednar, A.K., Cunningham, D., Duffy, T.M., Perry, J.D.: Theory into practice: How do we link? In: Anglin, G. (ed.) *Instructional Technology: Past, Present and Future*, Libraries Unlimited, Inc., Englewood (1991)
3. Davison, J., Arthur, J.: *Active Citizenship and the Development of Social Literacy: a case for experiential learning*. Citizenship and Teacher Education, Canterbury (2003)
4. Gratch, J.: mile: Marshalling Passions in Training and Education. In: 4th International Conference on Autonomous Agents, ACM Press, New York (2000)
5. Gratch, J., Marsella, S.: Tears and fears: Modeling emotions and emotional behaviors in synthetic agents. In: *The Fifth International Conference on Autonomous Agents* (2001)
6. Henriksen: On the Transmutation of Educational Role-Play: a critical reframing to the Role-Play to Meet the Educational Demands. In: Montola, M., Stenros, J. (eds.) *Beyond Role and Play - Tools, Toys and Theory for Harnessing the Imagination*, pp. 107–130 (2004)
7. Lewin, K.: *Field Theory in Social Science*. Harper and Row, New York (1951)
8. Marsella, S., Johnson, W.L., LaBore, C.: Interactive Pedagogical Drama for Health Interventions. In: 11th International Conference on Artificial Intelligence in Education, Australia (2003)
9. Mateas, M., Stern, A.: Architecture, authorial idioms and early observations of the interactive drama Façade. Technical report, Carnegie Mellon University (2002)
10. Mateas, M., Stern, A.: FAÇaDE: An experiment in building a fully-realized interactive drama. In: *Game Developer's Conference: Game Design Track* (2003)
11. Mateas, M.: Towards integrating plot and character for interactive drama. In: *Socially Intelligent Agents: The Human in the Loop*, AAAI symposium (2000)
12. Ortony, A., Clore, G., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge University Press, Cambridge (1988)
13. Magerko, B.: Evaluating Preemptive Story Direction in the Interactive Drama Architecture. *Journal of Game Development* 2(3) (2007)
14. Piaget, J.: *The Principles of Genetic Epistemology*. Routledge & Keegan Paul Ltd., London (1972)
15. Dias, J.: FearNot!: Creating Emotional Autonomous Synthetic Characters for Emphatic Interactions. MSc Thesis, Instituto Superior Técnico (2005)
16. Dias, J., Paiva, A.: Feeling and Reasoning: A Computational Model for Emotional Characters. In: Bento, C., Cardoso, A., Dias, G. (eds.) *EPIA 2005*. LNCS (LNAI), vol. 3808, pp. 127–140. Springer, Heidelberg (2005)
17. Weyhrauch, P.: *Guiding Interactive Drama*. Ph.D. diss., Dept. of Computer Science, Carnegie Mellon (1997)
18. Saretto, C.J., Young, M.: Mediation in mimesis Liquid Narratives. In: *ACM Southeast Conference* (2000)

19. Young, M., Riedl, M., Branly, M., Martin, R.J., Saretto, C.J.: An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development* 1 (2004)
20. Michael Young, R., Riedl, O.M.: Integrating plan-based behavior generation with game environments. In: ACE2005: Proceedings of the 2005 ACM SIGCHI International Conference on Advances in computer entertainment technology, p. 370. ACM Press, New York (2005)
21. Harris, J., Young, R.M.: Proactive Mediation in Plan-Based Narrative Environments. In: Panayiotopoulos, T., Gratch, J., Aylett, R.S., Ballin, D., Olivier, P., Rist, T. (eds.) IVA 2005. LNCS, vol. 3661, pp. 292–304. Springer, Heidelberg (2005)
22. Charles, F., Cavazza, M., Mead, J.: Character-Based Interactive Storytelling. *IEEE Intelligent Systems* 4(4), 17–24
23. Charles, F., Cavazza, M.: Exploring the Scalability of Character-Based Storytelling. In: Proceedings of AAMAS 2004, vol. 02, pp. 872–879. Springer, Heidelberg (2004)
24. Brisson, A., Paiva, A.: Are we Telling the Same Story? In: AAAI Fall Symposium on Narrative Intelligence Technologies, Westin Arlington Gateway, Arlington, Virginia, November 9–11 (2007)
25. Brisson, A., Dias, J., Paiva, A.: From Chinese Shadows to Interactive Shadows: Building a storytelling application with autonomous shadows. In: Workshop on Agent-Based Systems for Human Learning and Entertainment (ABSHLE) AAMAS. ACM, New York (2007)
26. Murray, J.: *Hamlet on the Holodeck: The future of narrative in cyberspace*, pp. 117–119. MIT Press, Cambridge (1998)
27. Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M., Woods, S., Hall, L.: Unscripted Narrative for Affectively Driven Characters. In: *IEEE Computer Graphics and Applications*, vol. 26(3), pp. 41–52. IEEE Computer Society, Los Alamitos (2006)

Making Stories Player-Specific: Delayed Authoring in Interactive Storytelling

David Thue¹, Vadim Bulitko¹, and Marcia Spetch²

¹ Departments of Computing Science
University of Alberta
Edmonton, Alberta, Canada, T6G 2E8

² Psychology
University of Alberta
Edmonton, Alberta, Canada, T6G 2E8
{dthue, bulitko, mspetch}@ualberta.ca

Abstract. Of all forms of storytelling, interactive storytelling presents authors with a unique opportunity: while most traditional stories must rely on having general high appeal, the nature of interactive stories to encourage audience interaction allows aspects of each individual's state to be automatically inferred. Given such information, an author's decisions would become more informed, and his ability to affect the audience would be improved. In this paper, we present an analysis of the decision-making process in interactive storytelling, and construct a method for characterizing storytelling systems based on features of their design. We demonstrate our method by comparing four recently published systems, and review related literature on inferring player information. Finally, we present Delayed Authoring, a new perspective on the design of interactive storytelling systems which takes advantage of their opportunity to make stories player-specific.

1 Introduction

The primary goal of storytelling is to have some impact on an audience, be it emotional, educational, or both. The author's success in achieving his desired impact, however, can greatly depend on each viewer's personal state, formed from her prior life experience and current state of mind. In traditional forms of storytelling (e.g., books, movies), authors have no access to this information; they are forced to take a generalized approach, relying on an understanding of their audience as a whole to guide their writing decisions. Interactive storytelling, on the other hand, allows a more customized approach, as information about the viewer can be inferred via the mechanism that supports her interaction [1,2,3,4]. Using such information, the author's writing decisions would become more informed, allowing a story to be created for each viewer that would be tailored to her state and have the impact that the author desired. While it would likely be unfeasible for a human author to create a unique story for every potential viewer, the technology behind interactive digital storytelling presents a viable alternative approach: express the potential elements of a story in a generic form, and construct an automated system to determine their details dynamically, informed by both an author's knowledge of the creation of stories, and the viewer's current state as inferred from her interactions.

Such a system would effectively serve as a decision-making proxy for the authors of a story, ideally choosing between potential elements of story content just as the authors would have, if they had been given a model of their viewer’s state to guide their story decisions.

Although having access to information about the viewer *while* a story is being both created and told is a benefit that is unique to interactive storytelling, the questions of both how to learn this information and how it might be well used (within the context of an interactive story) have only recently begun to be explored [1][2][3][4][5][6][7]. Based on this observation, the goals of this paper are twofold: to stimulate further research in learning the current state of an interactive story’s viewer, and to promote a new perspective on the decision-making process of interactive story design. To meet these ends, the remainder of this paper is organized into the following three contributions: (i) we present an analysis of decision-making in interactive storytelling, and use it to construct a new method for characterizing interactive storytelling systems based on the decisions made during their design and operation; (ii) following a review of recent methods for inferring a viewer’s personal state, we demonstrate our characterization method on four recently published storytelling systems, highlighting the ways in which they may be extended to take even further advantage of learned viewer states; (iii) we present a new design perspective, named *Delayed Authoring*, whose intent is to recognize the opportunities that arise when story decisions are delayed right up to the moment at which their results are needed. Finally, we conclude the paper with brief suggestions for a set of future research, showcasing the potential benefits that can be obtained by making interactive stories highly specific to a single, current viewer.

2 Decision-Making in Interactive Storytelling

Before beginning our analysis of decision-making in interactive storytelling, we define the primary subject of the decisions in a story, namely, story events, as follows.

2.1 Story Events

Fundamentally, stories are a sequence of events, each of which involves some form of action. In journalism, it is common to report the occurrence of an event by answering each of six questions: “Who?”, “What?”, “When?”, “Where?”, “Why?”, and “How?”. Considering the types of facts that result from answering these questions, we propose that events can be well-described by six general properties, as given in Table 1.

Table 1. Six general properties of story events

Property	Description
<i>Idea</i>	A brief description of the action that occurs.
<i>Actors</i>	The people/creatures/forces that either perform some action or are acted upon.
<i>Time</i>	The time at which the action begins.
<i>Place</i>	The environment(s) in which the action occurs.
<i>Actions</i>	The changes that actors make to themselves, other actors, or their environment.
<i>Reasons</i>	The notions held by actors that prompt their actions.

For example, consider this event: “*John rescues his friend Fred from the burning building by carrying him outside.*” In this case, the *Idea* might be “Rescue” or “Rescue from Fire”. The *Actors* are John and Fred, the *Places* are the building and the area outside, and the *Action* is John carrying Fred outside. Although the *Time* of the event is not stated explicitly, one might assume that the *Time* is “now”. One might also assume that John’s *Reason* for rescuing Fred is based on a desire to save his life. Although the given example fits well with the intuitive notion of story events being the significant elements of a story’s plot, our definition is more general: a story event may be created from *any* action, regardless of its significance in terms of the plot. For example, one might say that the *Reason* for an apple falling from a tree is the failure of its stem to resist the pull of gravity. Due to this generality, the following analysis, which focuses primarily on decisions related to a story’s plot, can be readily applied to character behaviour, dialogue generation, or any other story decision.

2.2 Story Decisions and Design Decisions

At its core, creating a story is about making decisions. For each of the six properties of an event, an author must decide what its value should be; that is, they must make six *story decisions* for every story event, as shown in Table 2. While the six questions shown all refer to a particular story event, consider also asking a similar set of questions with reference to each story decision itself; this set of questions is shown in the “Design Decision” column of Table 2. Note that the question asking “*where was the decision made?*” has been intentionally left out, as it is not relevant to this analysis.

Table 2. At left: six decisions to make for every story event. At right: five design decisions to make for every story decision.

Property	Story Decision	Property	Design Decision
<i>Idea</i>	What should happen?	<i>Result</i>	What was decided?
<i>Actors</i>	Who should be involved?	<i>Chooser</i>	Who made the decision?
<i>Time</i>	When should it happen?	<i>Time</i>	When was the decision made?
<i>Place</i>	Where should it happen?		
<i>Actions</i>	How should it happen?	<i>Method</i>	How was the decision made?
<i>Reasons</i>	Why should the actors act?	<i>Justification</i>	Why was the decision made in that way?

Just as the answers to the six story decisions in Table 2 yielded the properties of a given story event, answering the given five design decisions yields the properties of each story decision: *Result*, *Chooser*, *Time*, *Method*, and *Justification* (see Table 3). These properties describe details of the design of a given story experience, granting a better understanding of the storytelling system being considered. When considered in future tense (e.g., “*who will make the decision?*”), the questions from which they arose can be thought of as decisions that are made during the design of a storytelling system; we refer to them as *design decisions* (in contrast to story decisions) for this reason.

Table 3. Five general properties of story decisions

Property	Description
<i>Result</i>	The result of the decision being considered.
<i>Chooser</i>	The party who made the decision - either the player or the author.
<i>Time</i>	The time at which the decision was made - either offline (before the story) or online (during the story).
<i>Method</i>	The mechanism used to make the decision - this may be author imagination, a particular (computer) algorithm, or an in-game player action.
<i>Justification</i>	The author’s or player’s reason for using the method that they chose.

3 Characterizing Storytelling Systems

While story decisions determine the content of a story, design decisions determine a process by which stories can be made. In this section, we combine these two types of decisions to construct a method for characterizing storytelling systems.

For each of the six properties of a given story event, consider the properties of the decision that determined its value: What was the *Result*? Who was the *Chooser*? At what *Time* was the decision made, by what *Method* was it made, and what was the *Chooser’s Justification* for making it in that way? Answering this set of questions give rise to Table 4, which contains sample values to aid in explaining how it should be read.

Consider the first row of values in Table 4 (*Idea*). Reading from left to right, we learn that the event concerns some kind of rescue, and that this decision was made offline by the author, drawing freely from his imagination. The second row (*Actors*) shows that John and Fred are the actors in this event, and that this author decision was made at run-time (i.e., online) using an algorithm to automatically select two actors based on the requirements that the actors be in the same area as the event and friends with one another. The third row (*Time*) shows that the rescue event occurred in Act 3 of its story, and that it happened at this time because the player (perhaps being unable to save Fred

Table 4. Sample values for the five properties of each story decision

	Result	Chooser	Time	Method	Justification
<i>Idea</i>	Rescue	Author	Offline	Imagination	Unrestricted
<i>Actors</i>	John & Fred	Author	Online	Friend Finder	Plot Consistency
<i>Time</i>	Act 3	Player	Online	Persuaded John	Event Trigger
<i>Place</i>	Burning Building	Author	Offline	Imagination	Unrestricted
<i>Actions</i>	J Carry F Outside	Author	Online	AI Planner	Goal Satisfaction
<i>Reasons</i>	Save Life	Author	Offline	Imagination	Unrestricted

herself) found John and persuaded him to help. The fourth row (*Place*) shows that the rescue happened inside a burning building, and that the author chose this location from his imagination, offline. The fifth row (*Actions*) shows that the rescue event consisted of John carrying Fred out of the building, and that the decision for this action to occur was made online by an algorithm designed to plan John's actions toward achieving a given goal. Finally the sixth row (*Reasons*) shows that John's motive for rescuing Fred was to save his life, and that this motive was chosen offline from the author's imagination.

By filling out this table with respect to a representative event from a given storytelling system (and providing additional details in the "Method" column), one can concisely summarize that system's operation, characterizing it in the process. When an example event (such as the rescue event of Table 4) is made common to multiple systems, one can also highlight the inter-system differences and similarities that exist; we take this approach later on in this paper.

4 Inferring Player States

The current state of mind of any particular viewer causes them to approach a given story from a unique personal perspective. While traditional forms of storytelling offer little to no access to this information until after every aspect of a story has been fixed, the interactive nature of interactive storytelling allows viewers to convey their state (whether consciously or not) through the actions that they take while interacting. Henceforth, we will refer to the viewers of an interactive story as *players*, to better capture the facts that: 1) they are expected to perform actions *while* a story unfolds, and 2) their actions are meant to have some effect on the system that facilitates their story experience.

Several interactive storytelling systems have begun to model aspects of their players' personal states. What might it be useful for such models to represent? That is, what information about individual players would it be helpful for authors to have? Within the domain of interactive storytelling, the majority of relevant research to-date has focused on inferring one or more of three aspects of a current player's state: her knowledge, her preferences, or her goals. We review each of these aspects in the following sections.

4.1 Inferring Player Knowledge

Having a representation of what the player knows about the current story is especially important when the story is set in an environment that affords a high degree of player exploration, for players might easily miss seeing an element of content that is crucial toward understanding the circumstances of a subsequent story event. If this lack of player knowledge were correctly inferred, steps could be taken to either encourage the player to notice the content that she missed (perhaps via auditory or visual cues), modify the upcoming dependant event to rely on an element of player-known content instead, or switch to a different element entirely. Magerko's Interactive Drama Architecture (IDA) is an example of this approach, inferring its player's knowledge of the stories it presents by automatically tracking her movements between rooms in a virtual world [5].

4.2 Inferring Player Preferences

When an interactive storytelling system is aimed at providing entertainment, having some notion of the current player's preferences can aid the creation of a satisfying playing experience; the variation of styles of play or even basic interests from one player to another are typically too great for a one-size-fits-all approach. The common solution to this problem in commercial video games is to include a wide variety of content that is designed to appeal to different groups of players. The way in which this content is presented, however, often requires every player to either (i) experience every element of content in a largely sequential fashion [8], or (ii) manually search through a massive environment full of content to find the elements that they prefer [9]. If information about the player's preferences were correctly inferred, well-suited elements of the story's content could be automatically selected or brought to the attention of the player. Sharma et al. focused on learning the interests of a current player [3], Barber et al. and El-Nasr explored the task of learning player personalities (i.e., preferences in behaviour) [12], and Thue et al. aimed to learn their players' preferred styles of play [4].

4.3 Inferring Player Goals

In spite of the unique advantages that interactive storytelling provides, the very notion of allowing interactivity at all still causes great discomfort among many traditional story authors; the introduction of a character who is beyond their control (namely, the player) is often viewed as an obstacle to overcome, rather than a source of information to enjoy. Toward alleviating the reluctance of such authors, inferring a player's goals and intentions offers clues as to what the player is likely to do next. Albrecht et al.'s work on keyhole plan recognition follows this approach [10].

5 Using Player States

In this section, we employ the characterization presented in Section 3 to describe four recent interactive storytelling systems, all of which take advantage of inferred player state to guide their storytelling decisions. We refer back to our example of the "Rescue from Fire" event in Section 2 throughout, toward highlighting the similarities and differences between two pairs of systems: *Façade* and *Mirage*, and *IDA* and *PaSSAGE*.

5.1 Façade

Table 5 shows the story decision properties for Mateas and Stern's *Façade* [11] for the "Rescue from Fire" event. Even though *Façade* has enjoyed much praise from both the academic and consumer communities alike, Table 5 shows that it takes advantage of inferred player state in only two respects: determining the times at which events occur and the actions that actors take. A technique similar to the variable preconditions in *IDA* (Section 5.3) could add variety to the actors and locations of *Façade*'s events (given additional content), and certain players' interest in *Façade* could be improved by selectively introducing a more varied set of topics to discuss with the story's actors.

Table 5. Story decision properties for Façade

	Chooser	Time	Method	Justification
<i>Idea</i>	Author	Offline	Imagination	No Restrictions
<i>Actors</i>	Author	Offline	Imagination	No Restrictions
<i>Time</i>	Player & Author	Online & Offline	Tension Arc / Player Interest - a recent, sharp increase to the tension arc may have caused a high tension event to be sequenced. Alternatively, the player may have examined a fire alarm on a nearby wall, prompting John to begin his rescue.	Follow Dramatic Principle / Respond to Player
<i>Place</i>	Author	Offline	Imagination	No Restrictions
<i>Actions</i>	Player & Author	Online & Offline	Interruptible Scripts - High-level actor behaviours are pre-determined, but are authored in a way that they can handle player interruptions and then resume the original behaviour.	Allow Player Interaction
<i>Reasons</i>	Author	Offline	Imagination	No Restrictions

5.2 Mirage

Table 6 shows the story decision properties for El-Nasr’s Mirage [2], with respect to the “Rescue from Fire” event. Although Mirage infers substantially more player information than Façade, it uses its information to make primarily the same decisions: the *Time* of its events and the *Actions* of its actors. The main difference between the two, however, is that Mirage’s actors choose between different tactics based on predicted player behaviour, while Façade’s characters only cope with interruptions before returning to the same tactic as before. Mirage might take better advantage of its inferred player information by choosing actors for events based on an estimate of their tactics’ success against the current player’s character.

Table 6. Story decision properties for Mirage

	Chooser	Time	Method	Justification
<i>Idea</i>	Author	Offline	Imagination	Use Dilemmas
<i>Actors</i>	Author	Offline	Imagination	No Restrictions
<i>Time</i>	Player & Author	Online & Offline	Actor Improvisation / Player Model - John’s tactic for rescuing Fred may have been to convince the player save to him, but after predicting that this goal would fail, John’s tactic would change to rescuing Fred himself.	Character Believability
<i>Place</i>	Author	Offline	Imagination	No Restrictions
<i>Actions</i>	Player & Author	Online	Goal-directed Behaviours - Actors continuously monitor the potential success their goals, and choose new behaviours if failure is predicted.	Character Believability
<i>Reasons</i>	Author	Offline	Imagination	No Restrictions

5.3 The Interactive Drama Architecture

Table 7 shows the story decision properties for Magerko’s Interactive Drama Architecture (IDA) [5] for the “Rescue from Fire” event. Although IDA does take advantage of player knowledge to make part of its experience player-specific (namely, the time and place of its events along with some of the actions and actors involved), the “Time” column in Table 7 clearly shows the ways in which it could be even further customized (marked by “Offline”). For example, in the IDA testbed *Haunt 2*, while the murderer’s *Reason* for having killed the player’s character is completely pre-determined, some players might find an alternative motive to be substantially more interesting and compelling. It may be possible to learn such interests using work similar to Sharma et al.’s [3], and use them to further guide the selection of story events.

Table 7. Story decision properties for the Interactive Drama Architecture

	Chooser	Time	Method	Justification
<i>Idea</i>	Author	Offline	Imagination	No Restrictions
<i>Actors</i>	Player & Author	Online & Offline	Variable Preconditions - The identity of the character in need of rescue could have been left undetermined by the author, but constrained to be a friend of John’s.	Promote Variable Content
<i>Time</i>	Player & Author	Online & Offline	Satisfied Preconditions - The preconditions for this event might have been for Fred to be inside the building, with John and the player standing outside, and the building being on fire.	Use Available Plot Point
<i>Place</i>	Player & Author	Online & Offline	Variable Preconditions - The location of the fire may have been left as a variable by the author, allowing IDA to select the building that Fred is in to satisfy one of the event’s preconditions.	Allow Plot Point Flexibility
<i>Actions</i>	Player & Author	Online & Offline	AI Goal Selection or Reactive/Pre-emptive Direction - pre-scripted actions occur unless the player causes (or is predicted to cause) a boundary problem, at which time reactive or pre-emptive direction occurs.	Preserve Story Coherence
<i>Reasons</i>	Author	Offline	Imagination	No Restrictions

5.4 PaSSAGE

Table 8 shows the story decision properties for Thue et al.’s PaSSAGE [4], with respect to the “Rescue from Fire” event. Similarly to IDA, PaSSAGE makes significant use of inferred player information, basing the *Time*, *Place*, and *Actors* of its events on its model of the player along with her position in the story’s world. PaSSAGE differs from IDA in that it additionally chooses what should happen (i.e., the *Ideas*) in its stories, similarly to Barber’s GADIN [1] and other generative systems. A way in which PaSSAGE can further inform its decision-making lies in the *Reasons* for its actors’ actions.

Table 8. Story decision properties for PaSSAGE

	Chooser	Time	Method	Justification
<i>Idea</i>	Player & Author	Online & Offline	Encounter Selection via Player Model - This event would have been selected over others due to the current values in the player model.	Cause Enjoyable Events
<i>Actors</i>	Player & Author	Online	Role Passing - The identities of the characters in this event would be determined dynamically at run-time based on proximity and relationship constraints.	Satisfy Role Passing Constraints
<i>Time</i>	Player & Author	Online & Offline	Triggers for Encounter Specification - The fire may have been authored to start only when the player approached a suitable location.	Ensure Event Visibility
<i>Place</i>	Player & Author	Online & Offline	Triggers for Encounter Specification - The location of the fire may have been specified as the next building approached by the player which had one actor outside and another inside.	Bring Interesting Events to the Player
<i>Actions</i>	Player & Author	Online & Offline	Encounter Refinement via Role Passing with Hinting - Actors satisfying the encounter's trigger conditions would assume the behaviours that were authored for this event, tailored to encourage the player's preferred styles of play.	Preserve System Flexibility
<i>Reasons</i>	Author	Offline	Imagination	No Restrictions

6 Delayed Authoring

Having presented our analysis of decision-making in interactive storytelling and reviewed the ways in which players' personal states can be learned while they play, we now present *Delayed Authoring*, a new perspective from which to approach the design of interactive storytelling systems.

In the field of interactive storytelling at present, many researchers believe that a degree of tension exists between two stereotypically opposing forces: player desire, and author intent. The argument behind this point of view is straightforward; players wish to act freely in a virtual world, while authors wish to tell the best story that they can without contending with the whims of an uncontrollable main character. In truth, however, the unique ability of interactive storytelling to take player interaction into account is what allows this tension to be reduced, via the mechanism of player modelling. Making effective use of this mechanism is the primary goal of Delayed Authoring, as guided by three key principles of interactive story design; these principles are described in the following three subsections, and summarized at the end of this section.

6.1 A Decision-Making Proxy

The heart of the player versus author divide lies in the desire for control over the story experience, and this control is manifested in the types of story decisions that each party wishes to make. In the extreme case, players wish to choose when and where they travel in the virtual world and which characters they interact with, and they generally wish for

their actions to have a significant influence over the story that unfolds. Authors, on the other hand, wish to choose all of the details that the players would choose, along with a myriad of others as well, including the behaviour of secondary characters, the phrasing of lines of dialogue, and more. Recall, however, the information on which such author-driven story decisions are traditionally based: the author's knowledge of the creation of stories, and his general understanding of his audience as a whole. If the author knew more about his audience, his story decisions would be more informed, allowing his story to more effectively have the impact that he intended. This fact motivates the acquisition of the audience's personal state, and we have already discussed methods by which a player's knowledge, preferences, and goals may be inferred from her actions in a virtual world. It also motivates the *use* of the audience's state when making story decisions, but information about this state is only available after an interactive storytelling system has already been designed, built, and distributed to its players. The solution is to build an Artificial Intelligence (AI) component into the design of an interactive storytelling system which automatically both infers aspects of the player's state as the story is presented, and decides on subsequent elements of story content just as the story's authors would have, had they been asked to write a story for the particular current player themselves. While many existing interactive storytelling systems include AI to make story decisions, viewing the AI component as a decision-making proxy for the author is the first principle of Delayed Authoring.

6.2 A Just-in-Time Approach

We stated in the previous section that the greater an author's knowledge of his audience, the more informed his decisions would be. At what point in time would his decisions be *most* informed? Given that potentially useful information about the player's state could become available at any time, it should be desirable to delay the making of each decision for the maximum amount of time possible, i.e., right until the moment at which its result is needed. This just-in-time approach to decision-making is the second principle of Delayed Authoring.

6.3 Recognizing Opportunities

Referring back to the "Chooser" and "Time" columns of Table 4, three combinations of values are possible: Author/Offline, Author/Online, and Player/Online (assuming that players have no effect on the system before the story begins). For every Author/Offline pair, one might ask the question as to why it is not Author/Online instead; that is, for every decision that the author makes offline, one can ask if it could be made online, potentially benefiting from inferred player information in the process. The third principle of Delayed Authoring is to ask this question for *every* story decision encountered during the authoring process. Given the unique advantage of interactive storytelling to inform story decisions with learned player information, this principle encourages authors to recognize such opportunities whenever they arise. Table 9 lists the three principles of Delayed Authoring in summarized form.

Table 9. The three principles of Delayed Authoring

<i>Principle #1:</i>	The AI component of an interactive storytelling system should be viewed as a decision-making proxy for the interactive story’s authors.
<i>Principle #2:</i>	Any story decision that is made online should be delayed for as long as possible, to maximize its chance of being informed by new player information.
<i>Principle #3:</i>	For every story decision that arises during the authoring process, one should ask if it could be better informed by inferred player information.

7 Caveats and Future Work

Although the content of this paper has been focused on highlighting the potential advantages of a player-specific approach to interactive storytelling, there remain several important obstacles to overcome. Perhaps the most daunting of these is the problem of obtaining a sufficient amount of content to effectively *use* inferred player states at all; it is certain that traditional means of authoring will be unable to meet the varied demands of a large and diverse group of players. It may be the case, however, that the mechanisms for player-specific adaptation that have arisen in recent years can help to solve this problem, for the constructs underlying their ability to adapt are effectively representations of story events in an abstract form. Applying these adaptation mechanisms can vary the details of such events combinatorially, potentially allowing the procedural generation of a large amount of content with comparatively little authoring effort.

In addition to the four interactive storytelling systems whose characterizations we compared, it would be interesting to investigate other systems in the same context as well, particularly other generative storytelling systems such as Barber et al.’s GADIN [1] or Riedl et al.’s Automated Story Director [2]. Furthermore, in addition to a story’s content itself, the expression of that content can be made player-specific as well (e.g., Mirage, Façade [2][1]); in general, the principles of Delayed Authoring are meant to apply to *every* decision in storytelling, whatever those decisions may be.

8 Conclusion

In this paper, we made the following contributions. First, we presented an analysis of the decision-making process in interactive storytelling, conceptually distinguishing between story decisions as the determiners of story events, and design decisions as the determiners of how story decisions are made. We used this analysis to present a method for characterizing storytelling systems with respect to their decision-making processes, and demonstrated its use by comparing four recent interactive storytelling systems. We provided a comprehensive review of the literature on inferring player state from a player’s actions during a story, distinguishing between the player’s knowledge, preferences, and goals as different types of information to infer. We then presented Delayed Authoring, a new perspective on the design of interactive storytelling systems devised to alleviate the tension between player desires and author intent, encouraging authors to recognize the many opportunities to make their stories player-specific.

Acknowledgements

This research was funded by grants from NSERC and iCORE. Parts of this text appear in revised form from David Thue's M.Sc. dissertation at the University of Alberta.

References

1. Barber, H., Kudenko, D.: Dynamic generation of dilemma-based interactive narratives. In: 3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2007), pp. 2–7. AAAI Press, Palo Alto (2007)
2. El-Nasr, M.S.: Interaction, narrative, and drama creating an adaptive interactive narrative using performance arts theories. In: *Interaction Studies*. 8(2), pp. 209–240. John Benjamins Publishing Company, Philadelphia (2007)
3. Sharma, M., Ontañón, S., Mehta, M., Ram, A.: Drama management evaluation for interactive fiction games. In: AAAI Fall Symposium 2007 on Intelligent Narrative Technologies. AAAI Press, Palo Alto (2007)
4. Thue, D., Bulitko, V., Spetch, M., Wasylishen, E.: Interactive storytelling: A player modelling approach. In: 3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2007), pp. 43–48. AAAI Press, Palo Alto (2007)
5. Magerko, B.: Player Modeling in the Interactive Drama Architecture. PhD thesis, Computer Science and Engineering. University of Michigan, Ann Arbor (2006)
6. Nakasone, A., Ishizuka, M.: ISRST: An interest based storytelling model using rhetorical relations. In: Hui, K.-c., Pan, Z., Chung, R.C.-k., Wang, C.C.L., Jin, X., Göbel, S., Li, E.C.-L. (eds.) *EDUTAINMENT 2007*. LNCS, vol. 4469, pp. 324–335. Springer, Heidelberg (2007)
7. Peinado, F., Gervás, P.: Transferring game mastering laws to interactive digital storytelling. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) *TIDSE 2004*. LNCS, vol. 3105, pp. 48–54. Springer, Heidelberg (2004)
8. Valve Corporation: *Half-life 2* (2004), <http://www.half-life2.com/>
9. Bethesda Softworks: *The Elder Scrolls IV: Oblivion* (2006), http://www.elderscrolls.com/games/oblivion/oblivion_overview.htm
10. Albrecht, D.W., Zukerman, I., Nicholson, A.: Bayesian models for keyhole plan recognition in an adventure game. In: *User Modeling and User-Adapted Interaction*, vol. 8(1-2), pp. 5–47. Kluwer Academic Publishers, Norwell (1998)
11. Mateas, M., Stern, A.: Procedural authorship: A case-study of the interactive drama *Façade*. In: *Digital Arts and Culture: Digital Experience: Design, Aesthetics, Practice (DAC 2005)*, Copenhagen (2005)
12. Riedl, M.O., Stern, A., Dini, D., Alderman, J.: Dynamic experience management in virtual worlds for entertainment, education, and training. In: Tianfield, H. (ed.) *International Transactions on Systems Science and Applications*, vol. 4(1), pp. 23–42. SIWN Press, Glasgow (2008)

Verbal Communication of Story Facilitators in Multi-player Role-Playing Games

Anders Tychsen¹, Thea Brolund², and Michael Hitchens³

¹ Center for Computer Games Research, IT University of Copenhagen,
Rued Langgaards Vej 7, 2300 Copenhagen S, Denmark
anty@itu.dk

² University of Technology Sydney, Broadway, NSW 2007, Australia
theabrolund@gmail.com

³ Department of Computer Science, Macquarie University,
E6A, 2109 North Ryde, Sydney, Australia
michaelh@comp.mq.edu.au

Abstract. Multi-player role-playing games form one of the key examples of interactive, emergent and collaborative storytelling systems available. These games and the collaborative stories that they create, are commonly facilitated by a specialized participant, the game master. In the current study, the verbal communication of game masters in a series of role-playing game sessions is categorized and analyzed depending on form and content, using protocol analysis, establishing a model for the verbal communication of game masters.

Keywords: Game Master, Role-Playing Game, emergent narrative.

1 Introduction

Role-Playing Games form one of the major genres of games, represented in non-digital, digital and embodied formats. Depending on the form, role-playing games (RPGs) can be played with a single to thousands of participants interacting simultaneously in physical as well as virtual environments. Several varieties represent examples of collaborative, emergent interactive storytelling systems [2]. This is especially true for the classical expression of RPGs, that of the table-top RPG where a handful of participants partake in a collaboratively generated story framed by rules. No digital systems exist that surpass the flexibility of tabletop RPGs - not even player-controlled modules run with e.g. the *AURORA* and *ELECTRON* game engines for the PC game *Neverwinter Nights* [14], and they therefore form a rare example of a very pure interactive storytelling system. Just like in digital interactive storytelling systems, the players - or users - take on the role of fictional characters, operating in a fictional environment [3]. Inherent in RPGs is the same key problem that are facing designers of character-based, emergent interactive storytelling systems, that of reconciling the need for a coherent plot structure with the freedom of the user to affect the unfolding narrative [9, 10]. Table-top RPGs solve this problem by dividing the authorial control of the collaborative story - there are at least two - unevenly between the participants. Typically one (in some cases more than one) of the participants take

on the mantle of managing the overall flow of the storyline by taking control of the fictional world the player-controlled characters (PCs) operate within and determining how the world and its inhabitants react to the actions of the PCs within it. This concept of story facilitators are common to all RPGs, even live action role-playing games and massively multi-player online RPGs [4], and the functionally is named differently in different games, with the most common term being Game Master (GM) [7]. The operations of GMs form a source of inspiration for the construction of digital storytelling systems, because their function in effect is to reconcile the demands of a pre-authored plot with user influence (as well as user motivation and interest with engaging in the storytelling activity in the first place) [2, 12], in the process utilizing NPCs (autonomous agents). Existing work on GM functionality has addressed the basic processes of non-digital RPGs [8, 14]; the question of division of authorial control, and story planning and updating in runtime [2, 9, 12]. However, while this work has taken initial steps in developing models for the operations of GMs, it does not address the questions of what GMs actually verbally communicate to the players (users) during a multi-player RPG session? [14] presented an initial study of verbal communication in multi-player RPGs across tabletop and digital platforms, however, the focus was on the communication between all the participants; the GM was not considered as a separate entity. In the current study, protocol analysis [6] was applied to transcribed utterances from RPG sessions, which are categorized depending on form and content, and additional context-analysis applied to form an initial model of the verbal communication of GMs.

2 Experimental Setup and Methodology

A series of five PnP game sessions were run with groups of five players and one GM in a laboratory featuring one-way mirrors to permit participant observation without intrusion. Each group played the same RPG module – a form of pre-planned overview of the game story which the GM utilizes to facilitate the collaborative story - and negotiated the premises for the division of authorial control individually. The sessions were recorded, covering approximately 28 hours of game play. The game module utilized in the experiment utilized simple game rules, comparable with the *Dungeons & Dragons* D20-system. The module featured a sci-fi themed story, which pitches the player character in the middle of a war against a race of - alleged - terrible aliens. Characters were pre-defined.

Verbal communication and participant behavior was subsequently transcribed for approximately 40-60 minutes of playing time from five PnP sessions, distributed over three scenes from the beginning, middle and end of each game session, each with varying game story content (to accommodate variations in communications, as a function of playing time and game-story content). The transcriptions were subsequently subjected to protocol analysis, a method utilized within communications for studying both groups and individuals [6]; and applied in the context of multi-player games by [13,14]. Protocols are in this context recorded behaviour of the study objects, represented by video and audio recordings [1]. Protocol analysis generally produces large sample sizes and numerical data for which standard variations and similar statistical measures can be calculated. The data were analyzed in terms of raw

numbers of different types of categories, based on the concept of **utterances**, defined by [8] to consist of: 1) A subject who performs the communication 2) The content of the communication and 3) An object or objects to whom the communication is addressed. These properties were recorded for each utterance. The coding scheme used was originally developed by [15], and later modified for the current and other studies, using a combination of RPG theory to develop an initial model, followed by iterative development by test-coding digital and non-digital RPG sessions. The use of utterances as the basic unit of the coding hierarchy ensures a measure of degree and precision, however, the context of the communication may be lost. Therefore, more context-sensitive analysis was performed to find patterns in GM communication.

3 Results

The results were averaged across five PnP game sessions. The verbal communication of the GMs comprised a substantial 31.63% of the total communication among the six game participants in each group (standard deviation = 8.05%, range: 24.2%-45.05%).

GMs appear to communicate to both the entire group of players and one player only with approximately equal intensity: 43.79% (standard variation = 10.65%) of the total number of GM utterances were aimed at the whole group of players, with 56.21% (standard variation = 10.65%) directed at one player only. As indicated by the standard deviations, this was a highly variable pattern across the five PnP sessions examined. In rare cases the GM communicated with more than one player but less than the entire group. This typically happened when the fictional characters controlled by the players were separated within the game world geography. Of the total amount of verbal communication of the GMs, 96.8% was communicated out-of-character, with only 3.20% in-character (e.g. the GM role-playing a NPC) (standard deviation 1.70% in both cases). In the following example the GM is impersonating an army lieutenant: *“Okay everyone, I’m Lieutenant Sam Stone and we’ve got to go, this whole area is mined, the whole bunker could go up any second now.”* In the following, the GM is role-playing a nameless marine: *“Sir. The whole area’s mined, there’re grenades everywhere. We’ve got to get the civilians out of here. We’ll send the special bomb disposal units in to search for the weapons.”*

The low percentage of In-Character utterances indicates that the majority of the communication of the GM is focused elsewhere than in role-playing NPCs; however, this property of an RPG story is to some extent influenced by the type of RPG module being played. Some modules focus on the player-characters; others involve galleries of NPCs which play active roles in the game story. Furthermore, GMs have different preferences when it comes to communicating in-character vs. out-of-character. The low amount of In-Character utterances recorded in the recorded game sessions can therefore not be taken as indicative of RPGs in general.

These quotes are however evidence of another typical aspect of verbal communication in RPGs; that of using some degree of **dramatic language**, often combined with gestures such as the [salute] in the first example above. When communicating e.g. the behavior of an opponent or an NPC, a GM can utilize either “functional” language, i.e. describing the behavior without any attempt at dramatizing, or conversely try to use imagery or similar forms of dramatic emphasis,

to add dynamics to the situation. A functional statement could be: *“The bodyguard draws his sword, attacks you, and causes 14 points of damage”*. A more dramatically embellished version of this could be: *“The angered bodyguard draws his sword, roaring in fury as he takes a mighty swipe at you, cutting cleanly through your armor and giving you 14 points of damage”*. In the experiments, 36.90% (standard deviation = 6.25%) of the utterances of the GM’s carried some kind of dramatic flourish, with 62.44% (standard deviation = 6.24%) being wholly functional. One percent were purely dramatic, and had no game mechanic-functional content (standard deviation = 0.66%). The level and form of dramatic flourish varied, from addition of imagery to dramatic descriptions including body language, gestures and even the GM moving around the room, depicting action events (only verbal flourish was coded).

The use of dramatic flourishes varied between GMs. In general, dramatic flourish was used by GMs (and players) to add drama to situations, e.g. combat and similar scenes or events where the emphasis is on action and fast behavior. In these situations there is ample opportunity for GMs to use dramatic flourish in describing events. However, drama was also observed in situations involving social communication between e.g. an NPC and a PC. In these cases, dramatic language was utilized in a more tension-creating way, supporting the higher degree of intimacy in this type of character-driven scene. The use of dramatic flourish among the players varied, but similar to GMs tended to be most common during action sequences. Three forms were located: **A) Entity action** (recall that Fulzans are the name for an alien race): *“You unload through the breached doors, and with [GM makes sound of Fulzan roaring] one last surge, he just [GM mimics Fulzan ripping down door with both claws] rips the door apart with his fists, you know, his open claws. [...] and with that, the others come rushing in through the broken door.”* **B) Object/Environment:** Conversely, dramatic flourish is rarely observed in descriptions of the general environment of the PCs, or in descriptions of entities and objects within the environment. These statements tend to be more functional in nature: *“You look a bit around and about five hundred meters away from where they were supposed to be, there is some sort of a metal trap door that you can open. It looks like an entrance into a sort of bunker system.”* **C) Environment active property:** It is the actions of entities and objects that carry the greatest potential for drama, not the passive behavior of the environment. This is a rule with exceptions, e.g. when GMs described the horrible conditions of the alien planet, emphasizing the acidic atmosphere and the dangerous conditions. However note that these again relate to active properties of the environment, and how this might affect the player-characters: *“[Shakes head] No. And since everything that you can see out into the horizon is an artillery barrage, and the ground is still shaking under you, anything that has been alive is not alive anymore.”* Drama typically adds tension to the gaming situation, which is typically observable in the body language of the participants (e.g. very animated or focused). Exceptions to these principles occurred in every game session, so these are not rules set in stone however they provide a good approximation of the kinds of situations where the use of dramatic technique in storytelling are useful to keep the audience, i.e. the players, engaged and on their toes. None of the GMs used dramatic flourish all the time – rather they applied this in situations where it has effect, such as during combat, and for limited periods of time. This prevents the players from becoming “drama tired” and is a principle known from e.g. cinematography.

Content of GM communication: The content coding hierarchy groups the verbal communication of GMs into a system of coding categories, which cover the different forms of communication occurring in multi-player RPG situations, including in-game and out-of-game communication. Verbal utterances were coded based on content. The coding scheme included 28 content code categories defined under five groups: Assistance, Critique, Actions, Ask for Info and Other (table with full hierarchy omitted for space reasons). Of the 28 content code categories, eight had a frequency of more than 2% of the total number of utterances by the GMs and are discussed here.

Describing the actions of NPCs, providing information about the game world, rules or similar; forms the three most common categories, each comprise more than 10% of the GM-based communication, and are here discussed in more detail. The GM also requests information from the players, notably about the actions of their characters (5.92%), with other types of information requests measured (rules, game world, environment) forming under 1% of the total communication. 4.5% of the utterances were coded as “social communication”. These did not appear to hold any directly game-functional value.

Character Action Descriptions (CADs): These describe the actions of a non-player character or player character, and average 11.74% of the GMs communication, varying between 5.8-19.84% of the total verbal communication (standard deviation = 5.36%). For example, in the following the GM is describing the actions of a group of – presumably evil – aliens called Fulzans in one of the RPG sessions: “*There are four large Fulzans, holding another Fulzan down behind a car and taking up protective positions.*” The GMs would often describe the actions of the player characters even if the players were in control of these and generally described the actions of their characters themselves. Generally, GMs could take a CAD by a player, and “replay” it, providing at the same time information about how the fictional game world responded. In the example above the player has just stated the PC that he controls is filming the Private (a soldier) putting on his armour. The GM confirms this action by repeating it, at the same time adding a game world response to the action of the PC. This is also a dramatic tool for the GM, e.g. if there is a need to add drama to a scene, the GM can “re-describe” the actions of a player character using more dramatic language. In stating CADs, GMs can even define actions of player character that the controlling players have not initiated. These will generally be non-intrusive, focused more on dramatic description or adding humour to an in-game situation. The balance between intrusive control of player characters and simply enhancing the game experience is however hair-thin.

Give information: The give information content code falls under the code category “Assistance” and covers both spontaneous and non-spontaneous (on request) provision of information. On average, 15.88% of the total verbal communication of the GMs focused on the provision of information, varying between 7.93-24.51% (standard deviation = 6.02%), indicating a substantial variation in the requirement for the GM to provide information across the five game sessions. The giving of information is generally linked with an information request, i.e. non-spontaneous, by one or more players and fall within four general categories (Table 1).

Table 1. Information requests by players

Information requested	Description/example
Game world state	A request for an update on the state of the game world. By far the most common information request, encompassing about 50% of all player-based requests for information.
Other character actions/knowledge	Typically a request for an update on the placement and status of PCs and NPCs relative to the PC of the asking player/players.
Player character state	A request for an update on an action a PC has initiated, e.g. progress in putting on a suit of armour or swimming a river.
Rules	Requests for information about rules or clarification of the same, e.g. what number that is needed to roll on a dice to accomplish a specific action.

Responses to requests for information about the game world either take the form of environment descriptions, however also information about the general workings and mechanics of the game world. For example: *Player: "Aren't our cameras now set up to transmit to the A.T.T.A.C.K satellite?"- GM: "Your cameras are set up to transmit to the A.T.T.A.C.K satellite."* In this example the player is unsure about a specific property of two objects within the game world, the camera and the satellite, exists, and is asking the GM for a ruling on whether this property exists within the game world fiction. These kinds of information requests are important in setting boundaries for the possible actions of the player characters and serve to ensure the internal consistency of the game world fiction – e.g., in a RPG set in a fantasy-themed world, players asking if they have a laser pistol in their backpack will likely receive a negative reply from the GM. Related to the “give information” content code is “suggest/request action”. This code covered when the GM was providing hints to the players, suggesting they take specific actions, i.e. a form of help. In sessions where the GM held a more tightly controlling role, this could even go as far as the GM requesting the players to take specific types of actions. While rare, GMs can resort to this option, e.g. when players are about to take action which would kill their player characters, or disrupt the overall plot of the game story in a way the GM did not want.

Also related to the internal information flow is “information request (action of character)”, which covers the typical situation where GMs ask the players what actions their characters are taking. This question was typical of combat situations, where the actions of the PCs and NPCs have to be described at a resolution of seconds or less, in terms of time flow within the fictional world. Also typically related to combat situations is the “Acknowledgement” code, which consists of the GM confirming actions of the player characters within the game world.

Environment description: The most commonly occurring content category code was “environment description”, encompassing 47.49% of the total verbal communication averaged across the five game sessions (varied from 40.2-54.3%, standard deviation = 5.78%). In traditional multi-player, non-digital RPG situations like the one analysed here, the GM has a large degree of authorial control of what happens within the fictional game world and how it responds to the actions of the player. Communicating what happens within the fictional world and how it responds to the actions of the player characters, and updating the state of the fictional world on a running basis as

the players' progress through the game story, is a key task of the GM [2,9]. Descriptions of the environment cover all forms of sensory input, from visuals, sound, taste and even emotional impact of events within the fictional game world.

Environment descriptions come in a variety of forms, generally either describing the general appearance of the environment and **objects** within it, which surrounds the player characters, or describing the actions by **entities** within the game world. E.g.: “[...] which is the size of a 50-story office building. It’s this huge military base that has landed on the planet and is used to do the invasion. The invasion has to come from a ship because the atmosphere outside is so acidic that if you took off your helmet you’d die instantly as your eyeballs melt!”. Another example: “There are explosions; Fulzans run everywhere, they scream and yell. There are a bunch of Fulzans that take defensive cover, they’ve got themselves laser rifles, they’re shooting back, there’s gun fire going over everyone’s head.”

These can be combined, and often are in descriptions that cover the physical environment and objects within it, as well as the behavior of active entities. For example: “As you go into the armory, you see a shadow block the corridor [...] you come round the corner, the armory is there, the shadow that blocks the light out of the other end. Some creature, six foot tall and so bulky that it takes up most of the corridor, is lumbering down it slowly. This creature, you’re happy that you can’t see clearly - because it, but what you can see - it’s the most ugly, horrible creature in existence!”

Description of the actions of entities includes the player characters, as mentioned above. E.g., in response to a player noting that her PC will begin to put on a suit of armor: “The armor takes a long time to get in because you’ve got to put every single heavy piece on, and they’re all big heavy pieces and you don’t yet have the hydraulics activated. You start putting on your armor, and the doors which closed shudder as something hits it [bangs fist on table].”

The most common environment description of the GM is the response to the action/-s of one or more player characters: The GM responds to the behavior of the characters and updates the state of the fictional game world. This is exemplified in the below example, where a player has announced that his character is sitting down with a couple of soldiers to have a chat about the war: “So you are sitting there talking to the army guys about how the quality of the army can be judged on how bad the food is. This must be a bloody awesome army, because the food is really bad. “

Rules & social: While less common than the three most typical content codes, comments on the game rules or instructions to the players, for example what they need to score on a die roll to have their characters perform a specific action, were observed at regular intervals (frequency of 3.73%). The frequency of this content code was higher in sessions with less experienced players, where GMs used more time to explain the game rules both before/during the game session.

4 Discussion and Conclusions

The operations of GMs in table-top RPGs are a source for the development of heuristics for digital interactive storytelling systems. However, these games are highly complex systems. Before principles of GM operations can be transferred to digital

domains, a substantial amount of empirical research is required, focusing not only on communication between game participants in RPGs but also the game process [14] and the cognitive processes of GMs during play [2]. Furthermore, the requirements of each participant – their interests and desires experience – need to be examined in a comprehensive model of RPG play. In the current study, an initial step has been taken towards one of these areas of RPG play: The verbal communication of human GMs in table-top RPGs. GMs must balance a substantial amount of variables in their verbal communication with the players, perhaps most importantly in striking a balance between functionality and drama. In the above, the relative frequency of these elements and their contents has been analysed, however, future work on the present dataset will aim at analysing how GMs weigh the requirements of functionality, drama and continuity, and how they communicate this balance in practice; furthermore how the principles of this practice can be transferred to digital systems.

References

1. Akin, Ö.: *Psychology of Architectural Design*. Pion, London (1986)
2. Aylett, R., Louchart, S., Tychsen, A., Hitchens, M., Figueiredo, R.: *Managing emergent character-based narrative*. In: *Proceedings of INTETAIN 2008, Playa del Carmen, Cancun (2008)*
3. Cavazza, M., Charles, F., Mead, S.J.: *Character based interactive storytelling*. *IEEE Intelligent Systems* 17, 17–22 (2002)
4. Combs, N.: *The Intelligence in the MMOG: From scripts to stories to directorial AI*. In: *Proceedings of the OP Conference, Copenhagen, Denmark (2004)*
5. Cross, N., Christiaans, H., Dorst, K. (eds.): *Analysing Design Activity*. John Wiley & Sons, Chichester (1996)
6. Ericsson, K.A., Simon, H.A.: *Protocol Analysis*. MIT Press, Cambridge (1993)
7. Fine, G.A.: *Shared Fantasy: Role Playing Games as Social Worlds*. University of Chicago Press, Chicago (2002)
8. Gabriel, G.C., Maher, M.L.: *Coding and modeling communication in architectural collaborative design*. *Automation in Construction* 11, 199–211 (2002)
9. Louchart, S., Aylett, R.: *Solving the narrative paradox in VEs – lessons from RPGs*. In: Rist, T., Aylett, R.S., Ballin, D., Rickel, J. (eds.) *IVA 2003. LNCS (LNAI)*, vol. 2792, pp. 244–248. Springer, Heidelberg (2003)
10. Louchart, S., Aylett, R.: *From synthetic characters to virtual actors*. In: *Proceedings of AIIDE 2007*, pp. 88–91. AAAI Press, Marina del Ray (2007)
11. Manninen, T.: *Towards Communicative, Collaborative and Constructive Multi-Player Games*. In: Mäyrä, F. (ed.) *Proceedings of CGDCC*, pp. 155–169 (2002)
12. Peinado, F., Gervás, P.: *Transferring Game Mastering Laws to Interactive Digital Storytelling*. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) *TIDSE 2004. LNCS*, vol. 3105, pp. 48–54. Springer, Heidelberg (2004)
13. Smith, J.H.: *Plans and Purposes: How Videogame Goals Shape Player Behavior*. PhD-thesis. Center for Computer Games Research. Copenhagen, IT University (2006)
14. Tychsen, A.: *Role Playing Games – Analysis across Two Media Platforms*. In: *Proceedings of IE 2006*, Sydney, Australia, pp. 75–82 (2006)
15. Tychsen, A., Smith, J.H., Hitchens, M., Tosca, S.: *Communication in Multi-Player Role Playing Games – The Effect of Medium*. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) *TIDSE 2006. LNCS*, vol. 4326, pp. 277–288. Springer, Heidelberg (2006)

Improvisation and Performance as Models for Interacting with Stories

Theresa Jean Tanenbaum and Karen Tanenbaum

School of Interactive Arts & Technology,
Simon Fraser University-Surrey
250-13450 102nd Ave, Surrey, BC, V3T 0A3
{tess.tanen@gmail.com, ktanenba@sfu.ca@sfu.ca}

Abstract. One common metaphor for Interactive Storytelling has been the notion of Interactive Dramas, in which players assume the first-person role of the main character in a digitally mediated narrative. In this paper we explore the model of improvisation as a means of understanding the relationship between the author/designer and the reader/player of such narratives. This model allows for a new formulation of the notion of agency, by shifting the concept of the reader from a player-centric model to a performer-centric model. We also show how we can conceive of interactions between performers and authors as being governed by the same rules that are in play between multiple performers in a piece of improvisational theatre. We connect this idea to a phenomenological theory of human computer interaction and cognition which foregrounds the role of communication and commitment between interactors.

Keywords: Agency, Improvisational Theatre, Drama, Performance, Interactive Narrative.

1 Introduction

One common metaphor for Interactive Storytelling has been the notion of Interactive Dramas, in which players assume the first-person role of the main character in a digitally mediated narrative. One of the earliest scholarly explorations of digitally mediated narratives is Laurel's *Computers as Theater*, in which she outlines a model for Interactive Narrative based on theatrical performance. [1] In this paper we explore the implications of using the model of improvisational theatre as a way to understand the relationship between the author/designer and the reader/player of such narratives. This model allows for a new formulation of the notion of agency, by shifting the concept of the reader from a player-centric model to a performer-centric model. We conceive of interactions between readers and authors as being a co-performance mediated by the computational system, and governed by some of the same rules that are in play between multiple performers in a piece of improvisational theatre. We connect this idea to a phenomenological theory of human computer interaction and cognition which foregrounds the role of communication and commitment between interactors.

2 Players and Performers

There are two distinct, and occasionally conflicting, notions about the interactor in Interactive Storytelling and Games research. The first is the notion of the interactor as a *Performer*. This view is evident in early theoretical works, as when Murray likened an interactor's participation in digital environments to improvisational theatre, folk dancing, and jazz in *Hamlet on the Holodeck*. [2] It is found in the work of Laurel and later in the work of Mateas, both of whom used theatre and drama as models for theories and prototypes of interactive stories that place the player in the first-person role of *performing* the main character [1, 3]. In this notion, there is an expectation that the interactor assume a role, engage in the narrative as a character, and *act* in what might be described as a *narratively salient* way. *Performers* in interactive drama are engaged in the act of constructing a story in collaboration with the designer of the system, and with any other performers, be they AI agents or other human participants.

The second is the notion of the interactor as *Player*. This framing of the role of the interactor figures prominently in the arguments surrounding ludology. Espen Aarseth writes:

“In the adventure games where there is a conflict between narrative and ludic aesthetics, it is typically the simulation that, on its own, allows actions that the story prohibits, or which make the story break down. Players exploit this to invent strategies that make a mockery of the author's intentions.”[4]

The idea of players “making a mockery” of an author's intentions is rooted in a common assumption about ludic play that presumes that players will place the pleasures associated with winning over the pleasure of story. An interactor behaving as a *Player* in this context is more interested in what he can **do to** the game world, rather than what he can **do with** the game world. In this notion, there is an expectation that the interactor will behave in a self-gratifying, pleasure-centric manner, subverting the story in his own quest to satisfy his desire for *Agency*.

2.1 Agency

Agency is a term that has been much used and abused in game studies. Murray describes it as one of the “core pleasures” of digital systems, defining it as the “satisfying power to take meaningful action and see the results of our decisions and choices”. [2] It is around this notion that much of the debate surrounding the tension between the actions of the player and the control of the author has centred. In contemporary game design, agency is often interpreted to mean giving users the power to act without limitation or restriction. This notion might be better described as *unrestricted* or *true agency*. One of the goals of interactive narrative research is the quest to facilitate this *true agency* in a computational narrative environment, without sacrificing narrative coherence or quality. As a result, the extent to which an interactive narrative facilitates user agency is often one of the primary criteria for evaluating its success or failure. For all that discussion of agency and control has become commonplace in the field, Murray maintained that agency was *not* in conflict with authorial control, saying:

“There is a distinction between playing a creative role within an authored environment and having authorship of the environment itself. Certainly interactors can create aspects of digital stories in all these formats, with the greatest degree of creative authorship being over those environments that reflect the least amount of prescribing. But interactors can only act within the possibilities that have been established by the writing and programming.”[2]

This suggests that the agency advocated by Murray is a *limited agency*, where the designer defines the parameters in which the user can take meaningful actions. One of the key ideas in the above quote is the idea that moving away from pre-scripted choices allows a greater degree of creative freedom. This idea has taken root in the Interactive Storytelling community, manifesting as systems designed to adapt and respond to user actions procedurally [5-10].

It can safely be said that the range of possible interactions with Interactive Narrative and Games is broad enough to encompass both of these approaches. We argue that there is equal pleasure to be had in either the performance of a role in an interactive story, or in the expression of agency in a virtual environment. These two modes of behaviour are not mutually exclusive, and most players of games oscillate between them as they play. That being said, evidence in both scholarly research and games would suggest that most designers create systems with the expectation that interactors will approach them from the perspective of *Players*, rather than *Performers*. There is an assumption that most players want unrestricted agency, and are willing to sacrifice narrative coherence in order to get it. The common corollary assumption about the player is that if given a choice between free exploration and structured (narratively salient) performance, she will inevitably tend toward free exploration. Many games are even designed to reinforce this notion, by rewarding actions such as breaking open every crate or destructible object in a room to get power-ups and other ludic incentives. As a result, we have created a *culture of players* in games, which makes constructing interactive dramas that rely on *performers* difficult. One solution that commonly occurs in Interactive Narrative systems is to develop a guidance strategy which finds a way to creatively limit the player’s agency, while attempting to preserve the illusion of free will.

2.2 Expectations and Guidance Strategies

Guidance strategies manifest as the “rails” upon which players are steered in order to advance the story and are often sophisticated means of limiting the very player agency that designers assume will generate problematic situations in the narrative in the first place. Different theorists have proposed various approaches to this problem. Crawford argues that in order to balance the player’s inclination to act without narrative intent and the designer’s desire to tell a coherent story, the possible choices available to the player should be limited to only those choices which are narratively salient. He writes: “The storybuilder’s most important task is creating and harmonizing a large set of dramatically significant, closely balanced choices for the player.”[11] Other designers include various narrative “tricks” that enforce the author’s intent without violating the rules of the fictive world or directly denying the player action. Riedl, Saretto, and Young describe this as a tension between *control* and *coherence*. Their *Mimesis* architecture uses a technique they describe as

Narrative Mediation to respond to unanticipated user activity in two possible ways: *accommodation* or *intervention*. *Accommodation* allows the user to act freely, and attempts to “re-plan” the narrative structure around the unexpected action. *Intervention* alters the user’s action by “surreptitiously substituting an alternate set of effects”. The example they provide is of a user placing a coin in a Soda Machine and ordering a soda. The coin is necessary for a later piece of the narrative, and so the machine refunds the coin and declares the soda to be “out of stock”. [12] This is a form of *soft guidance*, in that it occurs within the fictive reality of the virtual world.

An alternative approach is to employ *hard guidance*. In the Soda Machine example, a hard guidance tactic might just prevent the player from interacting with the machine at all, perhaps including a piece of internal monologue from the character, such as “Nah, I’m not thirsty right now.” Regardless of which form of guidance is used, the result is the same: limitation of player agency. The hard guidance example that we have given, however, goes even further in that it denies the possibility of the player and the character occupying the same conceptual space. By suggesting that the player’s actions and the desires of the character are not in agreement, it enforces the notion that the player is merely *manipulating* an avatar, rather than *performing* a role.

The tendency to treat the interactor as a *Player* rather than a performer is quite common, especially in videogames, where play is the dominant interaction paradigm, but we contend that this is as much a result of the expectations and intentions of the *designers* of interactive experiences as it is the desires of *players*. While there is nothing inherently wrong with play as an interaction paradigm, it does place digital storytellers in a position where they must design *in spite* of the actions of their audience, rather than *in harmony* with them. In order to design an Interactive Narrative experience in which the goals of the player and the goals of the author are in harmony, it is useful to examine this other conception of the interactor: the interactor as *Performer*. There is some precedent for this approach: one of the most successful systems to emerge from scholarly research into Interactive Narrative, *Façade*, is built around Mateas’s notion of Interactive Drama, in which the player assumes the role of one of the main characters in the story. [7] Mateas and Stern argue for the incorporation of natural language recognition into the interaction paradigm of games in order to diversify the palette of possible experiences available to players. Implicit in this argument is the notion of players being able to communicate with the system the way they would communicate with other players – a notion which is borne out in *Façade*, which rewards players for “being proactive and acting dramatically”. [7]

3 Communication and Computational Systems

Winograd and Flores, in their 1986 book *Computers & Cognition*, discuss a new way of understanding both human and artificial intelligence¹. Their central contention is that the goal of creating human-like artificial intelligence is hopelessly mired in a

¹ The book is quite wide-ranging and thus it is hard to do justice in such a small space to the ideas it presents or the conclusions it reaches. We hope here to present a few salient ideas from the work without doing it a disservice by oversimplifying too much.

rationalistic worldview that reduces the complexities of lived experience to a set of objects, rules, and situations to be formalized and represented symbolically. They argue that this is not how people actually experience or act in the world, and put forth a phenomenological, Heideggerian notion of *being-in-the-world*, in which the majority of human action happens in a state of *thrownness* which has no stable representation and is not amenable to logical reflection and symbolic manipulation.

“[Computers designed based on rationalistic misconceptions] are restricted to representing knowledge as the acquisition and manipulation of facts, and communication as the transferring of information...[However,] computers are not only designed in language but are themselves equipment for language. They will not just reflect our understanding of language, but will at the same time create new possibilities for the speaking and listening that we do—for creating ourselves in language”[13]

They propose an alternative conception of artificial intelligence in which we recognize that all computational systems can truly do is reflect the intelligence of the people who designed them. Thus, the true potential for computers lies not in the mimicry or duplication of human intelligence but in the facilitation of human-to-human communication. To understand human communication and to expand on the idea of “creating ourselves in language”, the authors turn to Speech Act Theory, which classifies “what people can do with language” by identifying the illocutionary point of speech acts, such as to make a request, to promise something, or to declare a fact. Accomplishing such a dialogue relies on the idea of *commitment*: in the sense that both participants commit to enter into the conversation in good faith and also in the sense that they agree not to renege on any commitments made during the conversation. Crucially, these are commitments which are not just verbal and mental, but rather commitments to actions, positions and beliefs. Winograd and Flores hold that the ability to negotiate and commit to action via language is a key characteristic of humanity, and one which we cannot replicate via computers.

“Once we recognize the machine as an intermediary, it becomes clear that the commitment inherent in the use of language is made by those who produce the system. In the absence of this perspective it becomes all too easy to make the dangerous mistake of interpreting the machine as making commitments, thereby concealing the source of the responsibility for what it does.” [13]

To bring this notion of commitment and communication into the realm of Interactive Storytelling, the first step is to stop thinking about *the player* engaging with *the system*, and instead to focus on how *the author* and *the reader* communicate with each other via a computational medium. In addition to being a philosophical point, this has very basic, practical implications for how we discuss computational systems. The common use of language like “the computer decides to...” or “the system responds...” obscures the crucial fact that the system is not truly deciding or responding to anything; it is simply acting out the decisions and responses coded into it by the author or designer of the system. This may seem like sophistry, and it could be argued that such statements are simply shorthand for the intended meaning of “the system responds as it has been programmed to...” Winograd and Flores, however, argue persuasively for the power of language to shape, at a very basic level, how we

understand and experience the world, in which case the persistent use of language that over-attributes intelligence and commitment to computational systems can lead to people acting as if these characteristics actually applied. By talking about Interactive Narratives as a way for authors and readers to tell stories together, we can come to a different understanding of how to design such experiences.

4 Improvisation and Performative Knowledge

To accomplish this goal, we propose that the notion of player as Performer be explored more rigorously. This requires an investigation of the frameworks that have arisen in the performing arts, especially where improvisation is concerned. As is discussed in [1] and [14], improvisational theatre can be used as a metaphor for what Interactive Narrative designers would like to one day accomplish. Drama theorists Lockford and Pelias provide a good overview of performers and improvisers in their paper on *Bodily Poeticizing in Theatrical Improvisation: a Typology of Performative Knowledge*.

“Improvisational moments are engaged through an ongoing process of negotiation and coordination, through a positioning and repositioning of performers and their characters...Adapting to emergent circumstances, these performers are called to be aware communicators who can draw upon their cognitive, affective, and intuitive abilities...in order to absorb interaction details, create characters, and establish relationships. Establishing a communicative connection they must listen to each other and adjust their thinking and behaviour accordingly. They must incorporate new information spontaneously while also keeping an eye on producing a coherent narrative.”[15]

This description of improvisation incorporates vocabulary which should be familiar to researchers of Interactive Narrative. The highly contingent and emergent human process of improvisation is one which is done well only after years of training. It is not something that we have a clear formula for, even in human-to-human communication. Improvisation is a dialogue in which both participants must actively work to support and challenge each other, and if we are to take it as a model for Interactive Narrative systems, it is necessary that we recognise exactly how difficult a task it is. In the remainder of this paper, we examine how we might take the first steps in thinking about the design of Interactive Narrative from a performative standpoint.

4.1 Five Types of Performative Activities

Lockford and Pelias present what they describe as a “typology of performative knowledge”. They identify five different types of knowledge that live performers draw on in improvisation: communication, playfulness, sedimentation, sensuality, and vulnerability. Lockford and Pelias consider knowledge to arise from the body as well as the mind, and this framework treats improvisation as both an aesthetic and an *epistemic* act. Their term for this process is *bodily poeticizing*, and they use it to encompass knowledge that is “intuitive, somatic, affective, and cognitive”. [15] In Table 1 we present the key questions that they ask around each of the five types of performative knowledge.

Table 1. A Typology of Performative Knowledge [15]

Questions Asked of Performers
Communication
“Are the actors engaged in an ongoing process of negotiating and coordinating their characters and themselves through interaction? Do the actors seem connected, listening to and incorporating what each other is saying? Are they adjusting their thinking and action according to what they are hearing? Are they producing a coherent story?”
Playfulness
“Are the actors open to possibilities? Are they functioning with spontaneity and imagination? Are they playing with language? Are they recognizing linguistic and social constraints? Are they working within the limits of the given circumstances? Have the actors moved beyond established patterns to the ‘intricacies’ of the scene?”
Sedimentation
“Are the actors relying upon lifetime structures of learning? Are they trusting their bodies, following their impulses, paying attention to what feels right? Have they become reflective about their hidden, tacit knowledge? Have they considered the degree to which their sedimented behaviours match those of their characters?”
Sensuality
“Are the actors’ senses alive, ready, actively engaged? Are the actors taking in what they need? Are the actors feeling with their bodies? Are they open to the pleasures of sensory response?”
Vulnerability
“Are the actors willing to put themselves at risk? Are they willing to make difficult situations work? When feeling vulnerable, do they have the ability to keep focus on what needs to be accomplished? Are the actors willing to trust one another?”

The questions asked in this model give rise to a number of related questions that can be asked of Interactive Narrative, since the concerns that Lockford and Pelias express around improvisation are also present in the discourse of Interactive Narrative studies. For example, the questions asked above about *communication* are similar to how the field discusses *interaction*. Crawford’s model of interaction treats the phenomenon as a “dialog between two or more participants in which they alternately listen, think, and speak”. [11] This type of formalized dialectical interaction is very similar to the improvisational process of making and receiving “offers” that occurs between two or more actors in a scene. It is also congruent with Winograd and Flores’s treatment of communication between the designers of intelligent systems and their users.[13]

Similarly, the questions surrounding *playfulness* relate to the notion of working within constraints and limitations. Lockford and Pelias elaborate:

“To accomplish this playfulness, performers depend in part upon their ability to work within the inexhaustibility of language and to welcome its slippery excess...They know that even when choices are made, they are not the only choices available...At the same time, improvised moments also have boundaries to which the performer must be sensitive. Performers must follow linguistic rules, recognize theatrical conventions, and enact the given circumstances.”[15]

This treatment of playfulness recalls Zimmerman's definition of play, which he describes as the "free space of movement with a more rigid structure". [16] Within improvisational theatre, a number of rules and formal systems exist to help structure the play between actors. These rules facilitate creativity, communication, and play within the scene rather than limiting it.

The questions surrounding *sedimentation* can help address the extent to which the performer's behaviours and the character's behaviours align. This is an important issue for Interactive Narrative as there is a delicate balance to be maintained between designing a character with a distinctive set of behaviours and personality traits and designing a character that gives players the opportunity to express themselves within the narrative. Pearce discusses how game characters have identities that are very different from the characters created in traditional linear narratives.

"In typical narrative texts, both literary and cinematic, characters are central to the conflict. You cannot really imagine a story without characters. In a game, on the other hand, it is quite possible, and often desirable, to have a narrative with no 'characters' whatsoever...Games tend to favor abstracted personas over 'developed' characters with clear personalities and motivations. More abstracted characters leave more room for the player, and are therefore better suited to support a play-centric model." [17]

Presenting a player with an incomplete character invites her to fill in the gaps and allows her to perform the identity of the character however she sees fit. While this is one way to get character and player behaviour to align, it can also leave the player feeling at sea, without anything to base actions and decisions on aside from ludic concerns. Thinking about sedimentation can help in two ways: first, at the start of a narrative, playable characters can be incomplete but still contain sedimented clues about the character's imagined history. Second, sedimentation can occur over the course of the narrative arc. As the reader shares experiences with the characters, she begins to accumulate reasons for emotional attachment to them and investment in their story; this accumulation is the narrative equivalent to the lived process of sedimentation.

The idea of *sensuality* is perhaps the most difficult one to map onto Interactive Narrative. Until recently, Interactive Narrative systems and games were limited in their capacity to engage the interactor at a bodily level, and while recent phenomena such as the Nintendo Wii and the Sony EyeToy begin to show the potential of bodily interactions, the vast majority of interactive stories operate via more constrained interfaces. While many modern games are visually and even aurally engaging, the other senses are decidedly unused and in most cases the physicality of the interaction is limited to the fine motor skills involved in typing and button pushing. The exploration of alternative sensory experiences is an open and important research question. Finally, there is a connection between the notion of *vulnerability* and the simple process of "losing" or "dying". When a player fails a challenge in a game in a particularly irrevocable way, there is a moment of breakage: time must reset, a savegame must be loaded, the character must be resurrected, and the player must begin the task again. In this moment there is the possibility for learning and growth, but there is also possibility for frustration that might lead to the player putting the game down entirely.

5 Implementing the Performative Approach

The ideas of Lockford and Pelias provide a roadmap to follow when thinking about performative interaction, and in the next section we discuss how their questions can be adapted to specifically deal with the design of Interactive Narrative experiences. We believe that designing with these questions in mind will have a couple key effects on the experience of narrative. One is that designers will have to think critically about what it means to have a player *perform* in the story and what might be entailed in training players to do so. Second, a performative approach allows us to reconsider the notion of agency. Considering Lockford & Pelias's questions from a perspective of Interactive Narrative research, they highlight an aspect of interaction that often goes overlooked. This is the notion of trust between actors (or interactors), or what Winograd and Flores term *commitment*. In improvisation there is an explicit *social contract* between the participants in a scene. Each actor knows that she is responsible for making the other actors in the scene look good, or not. In order to improvise freely she must trust her partners to not make her look bad, and she must endeavour not to damage them in turn. In a computationally mediated interaction, there is less of an understanding of this social contract but it does not make it any less necessary. Designers of Interactive Narratives need to be able to trust their readers to engage with the experience in a manner that is friendly to the creation of a shared story, and readers of Interactive Narratives need to trust the storyteller to provide them with opportunities to express themselves within the context of narratively appropriate actions. The lack of discussion of this contract between performers in the field has resulted in narratives and games which are designed to "correct" for "unwanted" interactor behaviours rather than ones that encourage interactors to learn to participate in narratively salient ways.

We discussed above a few of the connections between Lockford and Pelias's typology and the discourse surrounding Interactive Narrative. Much of their model remains inaccessible to designers of Interactive Narratives; it is rooted in bodily experiences that don't readily translate into Interactive Narrative and Games. Moreover, theories of improvisation between human participants are able to assume approximately similar capabilities between all participants, whereas Interactive Narrative designers must consider the capabilities of the human interactors versus the level of responsive interaction that can be programmed into the system. In Table 2 we propose a number of questions that might be asked of the participants in an Interactive Narrative system, should we chose to regard it as a performance co-produced by an author and an interactor, mediated by a computational system. Lockford and Pelias's model is designed to diagnose improvisational scenes in order to determine whether or not the performers are succeeding at achieving a *known* correct state. Where their questions are *prescriptive*, and lead to a desired ideal performance, our questions are *analytical*, and are intended to explore issues of performance and improvisation in Interactive Narrative systems.

Table 2. Performative Knowledge for Interactive Narrative

Author	Interactor
Communication	
Does the author communicate narrative information clearly to the interactor? Does the author provide the interactor incentive to behave in a meaningful way? Does the author provide tools for clear and consistent communication? Does the author adjust the story in response to the actions of the interactor?	Does the interactor communicate clearly and consistently? Does the interactor behave in a way that is meaningful? Does the interactor understand the impact of her behaviour on the story? Does the interactor feel any sense of ownership or responsibility to the story?
Playfulness	
Does the author provide a range of meaningful possibilities to the interactor? Is there a coherent framework to restrict interactor actions?	Does the interactor explore various approaches to different situations? Does the interactor test the limits of provided boundaries?
Sedimentation	
Does the author allow the system to learn about the interactor over time? Does the story develop emergent behaviours and characteristics as a result of its history?	Does the interactor grow attached to the narrative over time? Does the interactor's personal history affect her experience and behaviour in the story?
Sensuality	
What "senses" does the author have available to tell their story? What role do these senses play in the response to the interactor?	How many of the interactor's senses might be employed to engage with the story? Does interaction trigger visceral or instinctive responses in the interactor?
Vulnerability	
What happens when the story "breaks"? Can it break in interesting ways? Is there potential for unexpected behaviour? Does the author expect the interactor to attempt to break it? How does the story recover from unexpected occurrences?	Does the interactor attempt to break the story? Is breaking the story a way for the interactor to achieve strategic gain, or enhance the quality of the narrative? Does the interactor receive any incentive for <i>not</i> breaking the story, or for performing within the expectations of the author?

5.1 Training Players as Performers

The example of improvisational theatre provides designers and theorists of Interactive Narratives with a valuable perspective for thinking about the identity of the interactor and the role of the author. If we think about Interactive Narratives as *computationally mediated performances*, we can begin to consider how to build narrative experiences that support the interactor in her role as a performer. This perspective also demands that we think about the way we manage the expectations of the interactor at a conceptual level. As discussed earlier, many existing works sidestep this issue by assuming that interactors are *not* going to treat the experience as a performance. This

is the equivalent of expecting one of the performers in an improvised scene to run around knocking over the scenery and throwing things at the audience. While that one actor may have an enjoyable time, she is not engaging in the project of creating a story with other actors.

It is not enough to simply say that interactors are now performers; we must take a cue from the explicit contract between actors in improvisation and design experiences that facilitate and encourage interactive performance as its own form of play. This does not necessarily come at the cost of ludic play—one of the five types of performative knowledge is playfulness, after all—however, it does require a shift in the relationship between play and agency. It also requires a transformation of techniques from improvisation and performance into techniques that can function in a computationally mediated performance. In order to treat the interactor as a performer, it is necessary to recognize that performing in any context is a specialized skill that is learned over time, and not an intuitive ability that every interactor may draw upon. In fact, Winograd and Flores say that even within everyday conversation, which most humans are perfectly adept at, there is still an advantage to be gained from explicitly studying and understanding the mechanisms governing the interaction in order to gain “communicative competence”.

“Communicative competence means the capacity to express one’s intentions and take responsibilities in the networks of commitments that utterances and their interpretations bring to the world. In their day-to-day being, people are generally not aware of what they are doing...Consequently, there exists a domain for education in communicative competence: the fundamental relationships between language and successful action. People’s conscious knowledge of their participation in the network of commitment can be reinforced and developed, improving their capacity to act in the domain of language. [13]”

Within improvisational theatre this notion of communicative competence is critical to the construction of a scene; actors spend years learning to clearly communicate with each other. While we do not propose that interactors within an Interactive Narrative should spend years training in improvisational theatre before they are allowed to enjoy a good interactive story, we do believe that it would be valuable to incorporate certain ideas from improvisational theatre into the culture of interaction with digital stories and into the design of these same experiences. Certain notions from improvisational theatre have value for designers and participants in interactive narratives: the notion of giving and accepting improvisational offers, the notion of an explicit contract between the participants, and the idea that every participant is responsible for the ultimate quality of the scene. These can all be used to inform Interactive Narrative design.

In order to design an interactive story that emerges out the contributions of both the author and the interactor it is necessary to imagine new ways of training the interactor in how to perform within the system. To a certain extent, many games already do this via extensive training levels, or ongoing training throughout the game. In the same way that we do not expect players of games to arrive knowing the specifics of the control scheme, or the game mechanics, we should not expect them to know how to communicate narrative meaning within the story. Instead of expecting interactors to

perform in a way that instinctively expresses narrative desire, or attempting to infer narrative desire from any and all interactor choices, we should explore techniques for training interactors in the specifics of a given narrative. In order for player's actions within a game to be of narrative significance, there needs to be a clear schema in place that maps interactor choices to story meaning. This schema needs to be explicit; it needs to be clearly communicated to the interactor so that she may master the manipulation of it in the same way that she masters the navigation of the space. In other words, it is necessary for interactors to build a communicative competence in their interactions with digitally mediated narratives. This need for communicative competence is equally important in the authoring and design of these narrative experiences; if the author is unable to communicate clearly with the interactor via the medium of the system, then the system requires reconsideration.

5.2 Reframing Agency

Re-framing of the identity of both the author and the interactor as performers engaged in co-creating a story helps us take a new perspective on the tension between narrative coherence (often described in terms of authorial control) and player agency. This tension can be traced to an implied hierarchy in the relationship between the designer of the story and the interactor. There is a sense that when a designer asserts something about the narrative, this assertion must come at the expense of the freedom of the interactor to choose for herself how the narrative should play out. In improvisational theatre, on the other hand, both actors are making assertions about the story in the form of dramatic offers. It is considered a violation of the rules of improvisational theatre to deny a dramatic offer, thus preventing either participant from denying the actions of the other. For instance, if Alan says to Barry "I heard about your dog's death...I'm so sorry," it would be a violation of the rules, and of Alan's trust, if Barry were to reply: "What are you talking about? I don't have a dog." By denying Alan's offer, Barry has not only made his co-actor look bad, he has also broken the narrative momentum of the scene. This denial of an offer is the improvisational equivalent of the guidance strategies described above in which the player attempts to buy a soda from a vending machine, only to be told (in one way or another) that she cannot get a soda from the machine.

If we can conceive of the relationship between the author and the interactor as one of equal participants in an improvisation, then the issue of interactor *agency* becomes one of performer *responsibilities*. From this perspective, each participant has a responsibility to accept the dramatic offers given by the other, and each has equal responsibility to the ultimate meaning of the narrative. In order for this performative dialogue to work, the author must design the system to be able to co-perform with the interactor. In improvisational theatre it is each actor's responsibility to make the others look good for the audience. The unfortunate corollary to this is that any one actor can easily make the other performers look bad. The contract between actors in this situation recognises the fact that all of the performers have a responsibility to the overall quality of the experience. This is similar to the notion of shared responsibility to the quality of a conversation described by Chris Crawford in his discussion of interactivity. [11]

6 Conclusion

Presently, very few games succeed at engaging the player in interactions at the narrative level. The current paradigm of interactive experiences has an implicit bias towards ludic exploration, rather than dramatic performance. One consequence of this is that a disproportionate value has been placed on *player agency*, with little attention given to how agency operates in improvisational performance. Given that improvisation and drama are both commonly used as models for Interactive Narrative, it seems only reasonable that contemporary theory from theater and drama be explored for its relevance to the task of designing Interactive Narrative systems. To this end we have adapted a model from improvisational theater to demonstrate how performance might be used as a lens for the design of Interactive Dramas. Implementation of this new model requires several significant conceptual shifts away from current theory and practice surrounding games. First, it is necessary to design systems that can train and guide interactors in how to perform within them. Second, it is necessary to reframe *agency* as a shared property of all participants in an interactive drama.

Central to this argument is the notion of the system as an *intermediary* between the author and the interactor, rather than as an autonomous entity. This idea has its roots in foundational work on cognition and intelligent systems. By considering storytelling systems as *communication channels* between authors/designers and performers, rather than as artifacts, it becomes possible to consider interactive narrative as a co-performance. This is a significant conceptual shift away from current theories of interaction and narrative that hold them to be in tension with each other. Winograd and Flores hold that the choice of language used to describe a phenomenon has an impact on our understanding of that phenomenon. We contend that this is true for the theory surrounding interaction with digital stories. In this paper we argue for a shift in the language surrounding Interactive Narratives, and a corollary shift in the underlying conceptual frameworks that guide the efforts of designers and theorists. While many of the challenges facing the field remain technological ones, we see much to be gained by clarifying and explicating the conceptual territory from which these artifacts and experiences will emerge.

References

1. Laurel, B.: *Computers as Theatre*. Addison-Wesley Longman Publishing Co., Inc., Boston (1993)
2. Murray, J.: *Hamlet on the Holodeck: the future of narrative in cyberspace*. MIT Press, Cambridge (1997)
3. Mateas, M.: A Preliminary Poetics for Interactive Drama and Games. In: Wardrip-Fruin, N., Harrigan, P. (eds.) *First Person: New Media as Performance, Story, and Game*, vol. 1, pp. 19–33. MIT Press, Cambridge (2004)
4. Aarseth, E.J.: Genre Trouble: Narrativism and the Art of Simulation. In: Wardrip-Fruin, N., Harrigan, P. (eds.) *First Person: New Media As Story, Performance, and Game*, pp. 45–49. MIT Press, Cambridge (2004)

5. Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M.: FearNot! - an experiment in emergent narrative. In: Panayiotopoulos, T., Gratch, J., Aylett, R.S., Ballin, D., Olivier, P., Rist, T. (eds.) IVA 2005. LNCS (LNAI), vol. 3661, pp. 305–316. Springer, Heidelberg (2005)
6. Cavazza, M., Charles, F., Mead, S.J.: Planning Characters' Behavior in Interactive Storytelling. *The Journal of Visualization and Computer Animation* 13, 121–131 (2002)
7. Mateas, M., Stern, A.: Facade: An Experiment in Building a Fully-Realized Interactive Drama. In: Game Developers Conference, Game Design Track, San Jose, CA (2003)
8. Seif El-Nasr, M.: Interaction, Narrative, and Drama Creating an Adaptive Interactive Narrative using Performance Art Theories. *Interaction Studies* 8 (2007)
9. Szilas, N.: IDtension: A narrative engine for Interactive Drama. In: TIDSE 2003: Technologies for Interactive Digital Storytelling and Entertainment, Darmstadt, Germany (2003)
10. Thue, D., Bulitko, V., Spetch, M., Wasylishen, E.: Learning Player Preferences to Inform Delayed Authoring. In: AAAI Fall Symposium on Intelligent Narrative Technologies. AAAI Press, Arlington (2007)
11. Crawford, C.: Chris Crawford on Interactive Storytelling. New Riders Games, Berkely (2005)
12. Riedl, M., Saretto, C.J., Young, R.M.: Managing interaction between users and agents in a multi-agent storytelling environment. In: Proceedings of the second international joint conference on Autonomous agents and multiagent systems, pp. 741–748. ACM Press, Melbourne (2003)
13. Winograd, T., Flores, F.: *Understanding Computers and Cognition*. Ablex Publishing Corporation, Norwood (1986)
14. Mateas, M., Stern, A.: Interaction and Narrative. In: Zimmerman, E., Katie, S. (eds.) *The Game Design Reader: A Rules of Play Anthology*, vol. 1, pp. 642–669. MIT Press, Cambridge (2006)
15. Lockford, L., Pelias, R.J.: Bodily Poeticizing in Theatrical Improvisation: A Typology of Performative Knowledge. *Theatre Topics* 14, 431–443 (2004)
16. Zimmerman, E.: Narrative, Interactivity, Play, and Games: Four Naughty Concepts in Need of Discipline. In: Wardrip-Fruin, N., Harrigan, P. (eds.) *First Person: New Media as Performance, Story, and Game*, vol. 1, pp. 154–164. MIT Press, Cambridge (2004)
17. Pearce, C.: Towards a Game Theory of Game. In: Wardrip-Fruin, N., Harrigan, P. (eds.) *First Person: New Media as Story, Performance, and Game*, vol. 1, pp. 143–153. MIT Press, Cambridge (2004)

Let's Pretend I Had a Sword

Late Commitment in Emergent Narrative

Ivo Swartjes, Edze Kruijzinga, and Mariët Theune

Human Media Interaction group, University of Twente,
Enschede, The Netherlands

{swartjes,kruijzingae,theune}@cs.utwente.nl

Abstract. We describe how autonomous character agents that inhabit a story world can make (out-of-character) decisions about its details, filling in the story world as they go. We describe how we model these kind of late commitment decisions, and discuss how we use them to support action selection and to justify the adoption of character goals. Although a rigorous evaluation remains future work, we have implemented the approach presented here and have performed some exploratory testing.

1 Introduction

The Virtual Storyteller [1] is a story generation framework based on the concept of emergent narrative (i.e., narrative through unscripted characters [2]). The aim of the project is to investigate possibilities of and issues with emergent narrative, in order to get insight into the conditions necessary for stories to emerge. The characters are enacted by intelligent agents, the design of which is informed by improvisational theatre [3]. We work from the hypothesis that character agents can produce better stories if they are able to make decisions in-character (IC, i.e., what does my character want, feel and do?) as well as out-of-character (OOC, i.e., what would be good for the story development?). For their IC decisions, the character agents run cognitive processes like cognitive appraisal and action planning. Some first steps towards endowing virtual agents with OOC decision processes have been made in the field, e.g., the “double appraisal” approach of Louchart, where characters choose actions based on their IC motivations as well as on the OOC aim to emotionally impact other characters [4]. In this paper, we elaborate on another OOC decision process, which we have called *late commitment* [3]. This concept allows virtual agents to fill in the story world in line with what the story needs. Props are added, characters are deepened, relationships defined, and the backstories of the characters are filled in *during the simulation*, when this is useful for story progression, rather than determined beforehand.

¹ Late commitment should not be confused with the term *least commitment*, commonly used in the planning community. While both have a similar rationale (i.e., keeping options open by delaying choices as long as possible), least commitment refers to deferring decisions about variable bindings and step orderings in the search for a plan, whereas late commitment refers to deferring decisions about the story world setting in the course of the simulation.

2 Late Commitment

Emergent narrative has previously taken an authoring approach whereby fixed, predetermined story world settings and character personality definitions serve as the basis for narrative development. However, such a fixed story world setting already considerably constrains the possible course of events, taking away some of its generative potential. Indeed, this approach can even be problematic since it is unclear from an authoring perspective *how* exactly the story world setting constrains the course of events, making it difficult to decide what props, relationships and backstories to put in the initial setting.

The idea of late commitment is based on the observation that in improvisational theatre, actors do not work with a predetermined, agreed upon story world, but frame it as they go along, continually adding new information by the things they say and do, and accepting the information that others contribute. To illustrate, children do this in improvised play (sometimes even very explicit, when they say: 'let's pretend I was the mother, OK?' [5]). In improvisational theatre this is never explicit but always implicitly conveyed through IC communication ('hello, *daughter*'), because OOC communication would disrupt the performance and the audience's suspension of disbelief.

When virtual characters are given means to modify the virtual story world, they can make similar framing decisions, but unlike real actors, they can have explicit OOC communication with each other 'under the hood', which remains invisible within the story world. We model these decisions using *framing operators*, STRIPS-like operators that embody OOC 'let's pretend...' activities. The execution of a framing operator leads to the assertion of its effects, but must create the illusion that they have always been true within the story world.

A similar idea has been briefly explored from a plot-centric approach by Riedl and Young. Their Initial State Revision (ISR) planner [6] is a story planner with support for a partially indeterminate set of initial facts, aimed to increase the space of possible plans (and consequently, the space of possible stories).

3 The Use of Late Commitment in the Virtual Storyteller

We have implemented support for the selection and execution of framing operators in the Virtual Storyteller and are currently experimenting with the authoring of a story domain about pirates, making extensive use of late commitment. We currently use late commitment in two cognitive processes: (1) goal management and (2) action selection. Both processes make use of a partial order planner, modified to allow the use of framing operators. First, goal management has been extended to allow the Character Agents to justify adoption of new goals when they have no goals to pursue. Character agents can adopt goals when the preconditions of these goals are met. For instance, adopting a goal of plundering another ship might require that there is another ship in sight, and that the character adopting the goal is a pirate captain. If these preconditions do not apply, they can be given to the character's planner as OOC goals, and can be achieved

by framing operators that introduce a ship in sight, and endow its character with the role of pirate captain. Second, action selection has been extended to enable agents to create plans for their goals. If the captain has adopted a goal to find out whether the approaching ship is friend or foe, he can make a plan involving looking through binoculars. If the story world does not contain binoculars yet, they can be framed to be in the captain’s cabin, affording a plan for the captain to go to his cabin to get them.

Preliminary exploratory testing revealed several issues with planning and execution of framing operators, which led to modifications of the agent architecture. When constructing a partial order plan using framing operators, care is taken that no IC actions are selected to satisfy preconditions of framing operators, which would yield a problem with observed character motivation (e.g., a character trying to become captain because it needs binoculars). Furthermore, if the planner selects an action whose preconditions contradict the effects of a framing operator, it always has to be ordered *after* the framing operator, to maintain the illusion that the effects of the framing operator were already true before any step in the plan.

There are also several consistency issues involved with the execution of a framing operator. For instance, if the illusion is to be maintained that certain effects have always been the case, all characters must unconditionally accept all of the framing operator’s effects, and the characters must be able to believably refrain from responding emotionally to these effects. To this end, a framing operator selected by one of the characters is first proposed to all other characters. If any character cannot consistently accept the framing operator, it rejects the proposal and the framing operator will not be executed.

Example. We will briefly illustrate the use of framing operators for goal management and action selection using a simple example domain about pirates. In this domain, there are two pirates: Anne Bonney and Billy Bones. One of the character goals defined in this domain is a goal to kill one’s enemy. The preconditions for adopting this goal specify that one must hate the character one wants to kill, but we haven’t defined any of the characters to hate the other. Instead, we have created a **HateYou** framing operator, that can be used to actively justify the goal. The framing operator defines that characters can just happen to hate each other². Anne can use it to pretend that she has always hated Billy, so that she can adopt the goal. Based on this goal, she plans to stab Billy, but for this she needs to be carrying a rapier. Again, we have not specified this in the initial setting, but added a framing operator called **CarryRapier** to the domain, usable by Anne to pretend she happens to be carrying a rapier. The preconditions contextualize the operator by specifying that (in our domain) it is believable that a pirate might happen to be carrying a rapier. Furthermore, they ensure that inconsistent situations are avoided, e.g., the situation in which the rapier Anne pretends to be carrying was already located somewhere else. Its effects state that the pirate is indeed carrying the rapier.

² Artistically, this is quite shallow, but it suits the example.

The example illustrates that we can abstain from thinking about whether Anne should have a rapier at the start of the story, and whether she should hate Billy, because the decision whether it makes sense to incorporate these facts is now made within the cognitive processes of the agents. The concepts of carrying a rapier and hating people must of course still be authored, but whether or not they end up in the simulation now depends on the decisions the characters make.

4 Conclusions and Future Work

The technique of late commitment contributes to agents that can introduce story material intelligently, i.e., when useful for the development of a story. We have introduced the notion of *framing operator*, referring to an OOC activity by one of the character agents to fill in previously unspecified knowledge. The benefit of late commitment is that it enables specifying the story world in terms of how it could be rather than how it is, and as such takes away the responsibility of the author to predict which exact properties the story world needs to have for a particular course of events. Furthermore, it offers the characters more flexibility in their reasoning processes; the characters can to a certain extent *choose* to adopt goals and enable actions, by filling in the world around them.

We have implemented a model of the concept of late commitment and are using it to support action selection and goal adoption. We have addressed some consistency issues that were revealed by exploratory testing, but a more rigorous evaluation remains future work. Furthermore, we believe the use of late commitment is not limited to these two character processes. The next step is to investigate how other processes can benefit from its use. For instance, framing operators might be employed to cause desired perceptions (e.g., seeing a hidden door can be enabled by introducing a hidden door in sight of a character), make story world events possible, or facilitate desired emotional reactions.

References

1. Swartjes, I., Theune, M.: The virtual storyteller: Story generation by simulation. In: Proceedings of the 20th Belgian-Netherlands Conference on Artificial Intelligence (BNAIC) (2008)
2. Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M., Woods, S., Hall, L.: Unscripted narrative for affectively driven characters. *IEEE Computer Graphics and Applications* 26(4), 42–52 (2006)
3. Swartjes, I., Vromen, J.: Emergent story generation: Lessons from improvisational theater. In: Intelligent Narrative Technologies: Papers from the AAAI Fall Symposium. Number FS-07-05 in AAAI Fall Symposium Series (2007)
4. Louchart, S., Aylett, R.: From synthetic characters to virtual actors. In: Proceedings Third Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2007) (2007)
5. Sawyer, R.K.: Improvisation and narrative. *Narrative Inquiry* 12(2), 319–349 (2002)
6. Riedl, M.O., Young, R.M.: Open-world planning for story generation. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), pp. 1719–1720 (2005)

On the Use of Computational Models of Influence for Managing Interactive Virtual Experiences

David L. Roberts¹, Charles Isbell¹, Mark Riedl¹, Ian Bogost², and Merrick L. Furst³

¹ School of Interactive Computing, College of Computing, Georgia Institute of Technology
{robertsd, isbell, riedl}@cc.gatech.edu

² School of Literature, Communication, and Culture, Georgia Institute of Technology
ian.bogost@lcc.gatech.edu

³ College of Computing, Georgia Institute of Technology
merrick.furst@cc.gatech.edu

Abstract. We highlight some of the characteristics of existing technical approaches to the management of *interactive experiences* and motivate *computational models of influence*, a technique we are developing to aid drama managers in the persuasion of players to make decisions that are consistent with an author's goals. Many of the existing approaches to managing interactive experiences have focused on the physical manipulation of the environment, but we argue instead for the use of theories from social psychology and behavioral economics to affect the adoption of specific goals by the player.

1 Introduction

Researchers have investigated *interactive narrative environments* as an approach to meeting ever-increasing user expectations of engaging virtual experiences. A *drama manager* (DM) [1] is a specialized type of experience manager [2] that attempts to balance the competing goals of creating a dramatic experience for the user while affording the user the freedom to act however they wish at any time. Almost all work on drama managers so far focuses on managing the player's experience by directly manipulating the world or instructing NPCs to interact with the player. In contrast, we propose using influence as the basic tool available to the drama manager. Thus, the physical manipulation of the environment typical of approaches to drama management is replaced with verbal cues frequently used in sales settings—the DM is, in effect, selling the player on the idea of adopting the author's goals.

We propose *computational models of influence* (CI), a novel technique that goes beyond the current approaches of physical manipulation. CI is based on theories of influence and persuasion from social psychology and behavioral economics. A drama manager built upon CI will be able to convince players to set their own goals to be consistent with the author's goals without them feeling manipulated or coerced. In doing so, we can move experience management from one of reasoning about physical manipulation of the environment to one of reasoning about *selling* the player on a goal.

We have three main motivations for this work: 1) provide authors with tools to help them sell their goals to the player; 2) simplify the authoring process by providing a principled approach to concretely implementing abstract authorial goals; and 3) maintaining or increasing the player's sense of self-agency.

Using computational models of influence, authors can create situations where a player is sold on the idea of adopting a long term goal rather than forced to accomplish a short term task. Take, for example, the task of trying to persuade a player to mow her lawn. In the past, this type of goal would be brought about through the physical manipulation of the environment using a type of “lock-and-key” approach: progress in the experience is halted (the lock) until the player performs the desired task (the key). A lawnmower would be placed in view of the player and (metaphorical) walls would be constructed to prevent the player from performing other tasks of significance. If the player does not mow her lawn, eventually she will realize there isn’t much else to do. Once the task is completed, the walls are taken down and the player can proceed.

Now, suppose it is desirable for the player to mow her lawn on a regular basis, not just once. One option is to repeatedly use the same object appearance/boundary construction technique described above. While that may be successful, it is somewhat limiting and may become repetitive or mundane for the player. By constructing a social interaction in the environment based on a model of influence, the drama manager can effectively sell the player on the idea of adopting the goal of mowing her lawn regularly. In doing so, the need to repeatedly manipulate the environment is eliminated but the player continues to comply with the author’s goal.

Further, CI will ease the burden of development on the author. Originally, gameplay progress was motivated by authoring an exhaustive set of local triggers—an enormous burden on the author. More recently, drama managers have brought to bear the power of AI on this problem and, as a result, have allowed for authors to be just as expressive while reducing the effort needed to specify how the gameplay will advance. Using computational models of influence, the ability of authors to stimulate progress will increase further by opening avenues for them to motivate more complicated and lengthy interactions by the player with little effort.

Many interactive narrative environments use a drama manager and so a number of systems have been developed. There is a wealth of work on drama managers and interactive narrative environments in general. However, space precludes a comprehensive discussion. Space precludes a detailed discussion of related work. Instead, see [3] for a survey of early work and [4] for a survey of more recent work. *Façade* [5] is a notable exception to physical manipulation-based drama management because it is designed to select beats to effect change in the mental or emotional state of the player. Beat selection is based solely on authors annotations and not on any accepted theory of influence or persuasion. In contrast, drama managers using computational models of influence will be able to automatically effect change in the players mental state in service of a physical goal—something the beat-based drama manager does not.

2 Computational Influence

Our goal in developing computational models of influence is to: 1) benefit authors by providing tools designed to influence players to buy into the adoption of goals consistent with the author’s; 2) reduce the burden on authors by enabling them to specify goals abstractly, relying on the principles of CI to bridge the gap to a concrete implementation in the virtual environment; and 3) accomplish (1) and (2) without the player perceiving any decrease in (and preferably an increase in) self-agency. We turn to the theories of

influence and *persuasion* from social psychology and the theories of *behavioral economics*. While there are a number of takes on those theories (a complete discussion of which is well beyond the scope of this paper), we have based the work described in this paper on the theories discussed by Cialdini [6].

All species—including humans—have certain built in mechanical responses to specific stimuli. These responses are called *click*, *whirr* responses to represent the mechanical click of a recorded tape loading and then whirring as it is played. In animals, it is believed that these responses are instinctual and free from social context. In humans it is believed that these responses are developed from psychological principles or social stereotypes learned over time. In fact, they are thought to be coping mechanisms. We use them to reduce our cognitive burden when dealing with the ever-increasing complexity of stimuli we are regularly faced with.

In order to use these principles effectively for interactive experiences, we need only to hit upon the trigger features that cause humans to play their recorded tapes. We focus on six principles of influence. These principles have been identified by years of research in the field of social psychology and behavioral economics and are frequently employed as sales tactics by savvy marketers. They are:

- **Reciprocation:** give and take; when someone does something for us we feel obligated to return in kind.
- **Consistency:** we have a near obsessive desire to be (and appear) consistent with what we have already done or said.
- **Social Proof:** we look to others like us to determine the appropriate action to take.
- **Liking:** the more we like someone, the more likely we are to abide by her requests.
- **Authority:** we have a deep-seated sense of duty to authority.
- **Scarcity:** something that, on its own merits, holds little appeal to us will become decidedly more enticing if it will soon become unavailable to us.

The principles provide the foundation for understanding how to create the powerful tools used on a daily basis by sales professionals. We are developing computational implementations of those tools. While these principles can only increase the likelihood of compliance, we believe that it is an important feature of our approach that the player can always decide not to do what the DM is trying to get her to do. Thus, the affordance for self-agency is strictly preserved.

3 Concluding Thoughts

In this paper, we have argued for the use of computational models of influence as a tool to help drama managers exceed some of their current limitations. We have shown that models of computational influence have three important characteristics that contribute to this goal. First, when used appropriately, the models of computational influence upon which our drama manager is built can shape the experiences of players in novel ways. Further, the use of these models will help players to adopt long-term goals. This greatly increases the potential for drama managers to guide players through increasingly complicated experiences.

Second, the models we are developing alleviate the burden on authors of specifying implementation details for their drama managers. Rather than have the author specify how drama manager decisions map to a concrete instructions for reconfiguring the environment or for NPC interaction, the author specifies high level story goals. We anticipate that this will reduce the burden on the author significantly.

Third, to truly provide a player of an interactive experience with a sense of self-agency, it is necessary to endow her with the ability to define and pursue her own goals. A major benefit of our models of influence is they increase (or at the least preserve) the sense of self-agency the player experiences. The plotline moves forward by the player's choosing, not by forced action or physical manipulation. As a result, drama managers built upon the concepts of computational influence enjoy the benefit of increased power, reduced authorial burden, and an improved sense of self-agency for the player.

We expect to continue developing a demo to test ideas in computational influence. We are interested in studying the ability of authors (both with and without previous experience) to create experiences using our paradigm. We also seek to understand the player's perception of self-agency in the environment. Our evaluation will consist of participants using our authoring tools to create their own experience and also asking them to participate in an experience that we have authored.

Acknowledgements

This research was performed while on appointment as a U.S. Department of Homeland Security (DHS) Fellow under the DHS Scholarship and Fellowship Program, a program administered by the Oak Ridge Institute for Science and Education (ORISE) for DHS through an interagency agreement with the U.S Department of Energy (DOE). ORISE is managed by Oak Ridge Associated Universities under DOE contract number DE-AC05-06OR23100. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DHS, DOE, or ORISE. We also acknowledge the support of DARPA under contract No. HR0011-06-1-0021.

References

1. Laurel, B.: *Toward the Design of a Computer-Based Interactive Fantasy System*. PhD thesis, Drama department, Ohio State University (1986)
2. Riedl, M.O., Stern, A., Dini, D., Alderman, J.: Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4(2) (2008)
3. Mateas, M.: An Oz-centric review of interactive drama and believable agents. In: Veloso, M.M., Wooldridge, M.J. (eds.) *Artificial Intelligence Today*. LNCS (LNAI), vol. 1600, pp. 297–328. Springer, Heidelberg (1999); First appeared in 1997 as Technical Report CMU-CS-97-156, Computer Science Department, Carnegie Mellon University

4. Roberts, D.L., Isbell, C.L.: A survey and qualitative analysis of recent advances in drama management. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4(2) (2008)
5. Mateas, M., Stern, A.: Integrating plot, character, and natural language processing in the interactive drama Façade. In: *Proceedings of the 1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2003)*, Darmstadt, Germany (2003)
6. Cialdini, R.B.: *Influence: The Psychology of Persuasion*. Collins (1998)

Purposeful Authoring for Emergent Narrative

Sandy Louchart¹, Ivo Swartjes², Michael Kriegel¹, and Ruth Aylett¹

¹ MACS group, Heriot Watt University, Edinburgh, United Kingdom
{sandy,michael,ruth}@macs.hw.ac.uk

² Human Media Interaction group, University of Twente,
Enschede, The Netherlands
swartjes@ewi.utwente.nl

Abstract. Emergent narrative (EN) is a narrative concept in virtual reality that relies on emergence for a flexible shaping of stories as opposed to fixed pre-determined plots. This has consequences for the creative role of the author in an EN system. In this paper, we aim to clarify the actual function of the author by investigating what is exactly mediated in ENs and how this can be filled in by an author at design time.

1 Introduction

In 1999, Aylett introduced the concept of Emergent Narrative (EN) as a credible solution to the “narrative paradox” in virtual environments [1]. The narrative paradox, as a term, illustrates the contradictions between an autonomous user, free to move in a virtual world, and the desire to convey a satisfying coherent plot structure. The original concept has benefited from work carried out within the Interactive Storytelling (IS) community (e.g., [2,3,4]) and has reached a general consensus on its components, articulation and design. The development and implementation of the FearNot! application [2] has also contributed to disseminate the concept to a wider audience and highlighted practical authoring issues associated with its specific character-based design. While previous research focused on identifying the interactions between dramatic elements (i.e., characters, plot, events), users and authors [3], the experience of authoring for FearNot! combined with the IS community’s shift towards authoring, raises a number of fundamental questions about not only the form of EN but also about the *process* by which an EN system is created. In this paper, we are concerned with the actual shaping of an emergent narrative and we aim to identify the process by which an author can create and organize narrative content.

The nature of EN is such that the author can enjoy a certain freedom from general concerns on interaction, contextualization or continuity. However, the lack of apparent narrative structure raises issues about authoring responsibilities. While character interactions move an unfolding plot forward, these interactions have to be authored in the first place, leading to doubts as to the relationship between what the author intends, which interactions to author, and which narratives emerge as a result. This line between authoring, emergence and fore-sought plot lines is rather unclear and is dependent on one’s semantic interpretation of

emergence. The view – as rightly pointed out by Crawford [5] – that original stories would emerge from a system given that the system is complex enough is an unrealistic take on the issue. Our opinion is that emergence should not be associated with a lack of purpose in the authoring process. In this paper we argue for the consideration of purposeful authoring, by considering first of all the notions of authoring and emergence with respect to the process within which an EN is assembled and presented to an interactor (Sect. 2), and secondly by considering the creative ideas that an author is able to mediate using the concept of EN (Sect. 3). From these two considerations, in Sect. 4 we describe authoring implications from the side of both the EN process (through the metaphor of story landscape) and the EN as medium (by considering EN as simulation).

2 How an Emergent Narrative Takes Shape

In principle, an EN system is designed to offer a certain dramatic experience to an interactor. The interactor, by assuming a given role, takes on part of the responsibilities for the qualitative and interactive aspects of the experience. Such a dramatic experience can only take place if the interactor is actively participating and has been given the means to participate relatively freely with both the narrative environment and the characters that populate it. The EN concept approaches storytelling from a process-based perspective. In this particular context, rather than focusing on the structure of a given story, we propose to develop an understanding of how one should envision an emergent narrative to take shape as a process. We aim to carefully identify the roles that the interactor and system play in accomplishing a satisfying interactive narrative experience.

2.1 The Interactor: Narrative Development through Interaction

An important aspect of an EN system is that interactors influence how the narrative unfolds. In order to do so, interaction and narrative development must show a certain level of flexibility so as to accommodate each other. The shaping process between a story and an interactor’s choices reflects the decisions made by a user on the spectrum of actions or events proposed by an EN scenario. In this model, a given narrative system offers a definite range of options to an interactor. The narrative development is not pre-defined however; the decisions that the interactor makes gradually shape and re-shape the spectrum of actions available for a meaningful and purposeful experience.

This is illustrated in Laurel’s “flying wedge” (Fig. 1) [6]. An interactor within an EN determines the direction of the narrative development by engaging in certain interactions (the aim of the point of the wedge), and in turn the narrative development constrains the probable future interactions (the gradual narrowing of the wedge). In the case of EN, interactions are constrained by the narrative development through the interactor’s own storification process (i.e., the ongoing cognitive process of constructing a narrative understanding of the experience [7]). By storifying the unfolding sequence of events, the space within which future




Fig. 1. Laurel’s “flying wedge” (interactive version). In EN, interactors’ choices determine the direction of the narrative development and in turn, narrative development constrains the range of probable future interactions.

interactions make sense is reduced. At the beginning of an EN, anything is possible but by the time the first few interactions have taken place, this space is already considerably narrowed. The characters take their context along in deciding their future behaviour, and this context only keeps growing as it also contains their history up until that point. For instance, at the point where the wolf, in the Little Red Riding Hood folk tale, starts pretending to be the granny, interactions that would have been purposeful *before* the scene takes place, such as the wolf inviting granny for dinner at his house, make no sense from a dramatic perspective any more. The interactions that make sense from then on are the ones pertaining to the deception scene.

By making decisions as the story unfolds, the interactor determines the range of possible interactions that could logically or dramatically follow the course of action undertaken (illustrated by the direction the wedge points at). Therefore, as a result, certain actions would fall outside the scope of the interactor’s probable action range and would not be used.

Laurel’s Flying Wedge transposed to multiple characters means in EN that all characters have their own wedge. This illustrates how a particular range of actions can be utilized within an EN scenario. In this case, actions that would have been deemed unnecessary for a particular character might still be part of a potential range of actions for another character. While this hints at a definite story space for interactions to shape an EN, it does however not provide answers as to how this definite space should be created and elements within it authored. We will therefore elaborate on this discussion from an authoring perspective in Sect. [4.1](#).

2.2 The System: Virtual Characters and Drama Management

The prime concern in designing an EN is to assess the different natures of interactions and how they affect each of the characters’ stories. An EN approach facilitates such design issues by offering a virtual reality in which these interactions

can take place. One can rely on virtual characters to mediate the purpose of interactions. Since an EN story is conveyed through its characters, the interactions they engage in with each other and with the interactors should help the story to move forward. This should also be further supported by the dramatic modelling of characters as these are designed for *change* (in the sense that the interactions change the attitudes, emotions and goals of the characters).

An often mentioned issue with a character-centric approach is that interactions of autonomous characters do not necessarily lead to interesting stories. Although this might be particularly true when stories are considered from a spectator viewpoint (stories in the OZ experiment were not particularly exciting to watch [8]), this might be less so for stories experienced from a participative view. Still, as previously mentioned, emergence does not mean losing all concern with story development. It just means that the managing forces for drama in EN are necessarily operating in a way that is very opportune. Drama management in EN is not based on desired plots and an attempt to 'push' autonomous characters or interactors to do what is necessary, but rather uses a very local view on plot development where it should seize opportunities as they come along. For instance, drama management in FearNot! and the Virtual Storyteller [9] happens by system components that facilitate the story development by setting up episodes and scenes for the purpose of increasing the chances of story progression. Under certain conditions, lapses in time, space and situation can be made that will facilitate story progression by bringing parts of the simulation to the fore. More subtle forms of drama management are currently under development, where the agents acquire some responsibility to make decisions that aid on story development, such as choosing those actions that have the biggest emotional impact on other characters (the Double Appraisal mechanism [10]), or introduce events or new setting information to enable useful goals or actions (late commitment [11]). An implication of this opportune stance toward drama management is that the quality of the emergent stories might vary depending on the particular opportunities and the particular ways in which drama management decisions have played out.

3 Mediation in Emergent Narrative

It may be clear by now that EN is not meant to mediate stories as artefacts thought up in advance, because, as we have argued, the interactor has a fundamental role in 'what the story is', or rather, what *their* story is. This poses the question of what exactly EN conveys. If stories emerge at interaction time, one might wonder who "tells" these stories, and to whom they are told. Furthermore, what is the story being told? These questions pertain to a deeper underlying question about the EN concept: what can an author mediate using the concept of EN? We take a closer look at what mediation means in terms of EN, by investigating the roles of sender, receiver and mediated message.

3.1 The Sender: Who Tells?

In traditional narrative, authors take on full responsibility for how a story is received by the way they write it. When wanting to mediate a particular story to the receiver, they write a story in advance, skilfully raising questions the receiver might have, and answering them at just the right time to reach desired effects. This unidirectional relationship disappears in EN. Although an EN system is authored in advance, the responsibilities of narratorship become shared between system and interactor. It is the interactor who raises questions and sets out to answer them in the emergent narrative. As such, an EN is the story of the interactor, not in the sense that they have been the sole creative force behind it, but in the sense that it is in part driven by their attempt to organize their experiences into a unified whole. As such, an EN is established as a dialogue between system and interactor.

3.2 The Message: What Is Told?

Because an EN gives the interactor responsibilities that the narrator would otherwise have, the message or moral of an EN takes on a shape that is different from that of traditional stories. Where traditional stories bring across a unidirectional message, an EN gives the interactor the chance to construct their own message, by developing an understanding of the story world through making choices and exploring their consequences. For instance, *FearNot!* does not tell children how they should deal with bullying behaviour. Rather, it lets them try out for themselves what works and what does not. By authoring the various causes and consequences of a moral dilemma, the responsibility for the conclusion is placed in the shoes of the interactor storifying them, and as such becomes very personal. After many replays the conclusion might gain more nuance and sound something like: “often when I push the bully, he falls and stops bullying, but sometimes he doesn’t and the situation gets worse.”

3.3 The Receiver: Implications for Participation

Shared narratorship has far reaching implications on the interactor’s role within an EN and on the formulation of requirements for an author at design time. Without considering these, one might be tempted to think that an EN system should cater to a vast array of actions the user might want to do, and make sure they all have consequences in the story world. Therefore, we consider here the notion of agency (or meaningful action) [12] from the perspective of EN in order to relax this heavy responsibility of an EN system.

There are three points we would like to make about how we approach agency in EN. First, interactors do not have to be able to predict the consequences of their actions in terms of story outcome in order to be able to experience agency; understanding how particular actions effect the story line can even come much later. In an EN where the interactor plays Little Red Riding Hood, they might be persuaded by the wolf to take a detour, and unveil the wolf in granny’s

house. This might suggest to the interactor the reason *why* the wolf sent them on a detour. If this trickery was not predicted by the interactor, its realization contributes to their storification process, making sense of their own actions and those of the wolf only later, whilst still being satisfying. This implies that “purposeful action” does not necessarily mean that this purpose plays a central role in the story. This property sets it apart from most computer games.

A second point is that if the interactor *can* predict the consequences of a certain course of action, it should not be expected that they make the choices that they would if it were real life. For example, when the wolf addresses the interactor, they might remember the wise words of their mother telling them not to speak to strangers, but still decide to talk to the wolf so as to find out what happens (from a dramatic perspective). This implies that the interactions do not need to be tailored such that the interactor can “be themselves”; for a good story to emerge they might not only have to, but also *want* to make the dramatically bolder choices, just like the other characters should be doing.

The third point is related to the previous one, in the sense that EN presupposes a willingness to play within the formal constraints of a role. These formal constraints might be partially defined at the start of an EN but also establish themselves further during play in the form of *offers*. Through such offers, an EN establishes a certain cooperative contract with the interactor [13]. For instance, if the Little Red Riding Hood EN starts at home and mother calls the interactor into the kitchen, the offer intends for the interactor to go over and speak to her. They would not be going along with the offer if they ignored the call and went on their own little exploration. Our viewpoint is that the system is not responsible for keeping the interactor cooperative; we do not require it to cater for all ways in which they might not be going along (*blocking*). If the interactor blocks, the resulting EN experience might not be a very satisfying one, as user involvement is part of the fundamental basis of the concept.

4 Authoring an Emergent Narrative System

In the previous two sections we have described how an emergent narrative can be shaped at run time and what exactly this particular storytelling medium could mediate. However, how an author is supposed to think and work within an EN system is still an open question. We now refer back to those points from an authoring perspective. We first investigate through the metaphor of a story landscape how an author can create an interactive experience that unfolds as described in Sect. 2 in an attempt to get a grip of the space of possible stories the author is crafting out. Consequently, by discussing how EN is essentially a simulation approach to storytelling, we address the implications and pitfalls of authoring if the purpose is to mediate a story world.

4.1 The Story Landscape: Authoring for Narrative Development

In EN we try to remove the need to think in terms of plot, because the notion of plot – as discussed – has a problematic tension with the role of the interactor. To




Fig. 2. The story landscape with some of its many possible paths sketched

this end, an EN system models how characters come to make certain dramatic decisions, so that the author is then left to determine the content of the emergent narrative, which is raised to a more declarative level (which goals, actions and emotions there are, and under which conditions they occur). This way, the author can think directly in terms of these interactions and what happens locally, rather than in terms of the plots that should occur. This means that all the emotions, goals, thoughts and actions of the story domain are contextualized in terms of the character's autonomy, i.e., the author describes in which context they envisage that a character will feel a certain emotion, adopt a certain goal or perform a certain action. The advantage is that if this is done well, the behaviour of the characters is always in the right context.

To understand the emergence of plot based on authored material, we use the metaphor of a landscape of possible stories (see Fig. 2). Points on the landscape represent possible states of the EN, and climbing hills represents moving towards more and more dramatic necessity (cf. the Flying Wedge of Fig. 1). In a “valley”, there are many potential mountains to climb and many paths to do so. Character interactions move the dramatic situation more and more uphill since they yield emotions and intentions for the characters that form a reason for more specific behaviour. It is the author's task at design time to write behaviour that yields a story landscape for the interactor to travel upon, and it is therefore important to understand better the exact processes that constitute such a landscape.

The story landscape is a result of how local interactions of the characters play out and can therefore not be authored directly. EN being a character-centric approach to IS, the concrete content that the author provides is written from the perspective of a single character. The system has content knowledge of the different actions the characters can perform, the goals they can have, the emotions they might experience, and (either explicitly or implicitly) the context for the occurrence of these elements. The system uses this content knowledge together with procedural knowledge in the form of the cognitive processes of the characters (for instance, action planning, goal selection, cognitive appraisal and coping), to map out this landscape of possible stories.

Each character has its own landscape of possible stories. An EN system is aimed at meaningful interaction, which happens for instance when the actions of one

character establish the context for emotions or goals of other characters, whose performed actions again might lead to emotions and goals for a third. There is no steering force on how this plays out exactly, nor can the author envision this exactly, and it is this property – the real-time translation of autonomous action at character level to interaction at story level – that is emergent in an EN system. The landscape and its peaks are implied by the authored content through the way it is processed by the EN system. For instance, if Little Red Riding Hood takes on the goal to bring cookies to her grandma, this constrains her behaviour and – in terms of the metaphor – sets her on the way to a peak, which is a different peak from that in which her mother had asked her to wash the dishes.

While the “story landscape” viewpoint provides no obvious authoring solutions because the author creates this story landscape only indirectly, it does help in identifying certain authoring issues. In the following we focus on three authoring issues (boundaries, critical mass and dead ends), which can be identified by a critical observation of the story landscape, but none of which is inherent in a single item of authored content. We will provide some design suggestions on how content can be structured and shaped to tackle such issues.

Boundaries. A boundary is what separates the story landscape from the rest of the universe (the “sea” around the story landscape, if you will). For example there are no submarines in an EN about cavemen, because they fall outside the envisioned landscape of possible stories. An EN needs boundaries, not only because of the technical infeasibility of simulating an unconfined world but also because the boundaries help define the topic, scenario and message of the EN. This notion of boundary is however quite abstract and can be realized in many different ways. For example, one might construct spatial boundaries (given by the locations where the story takes place), contextual boundaries (e.g. the bullying context in FearNot!) and interaction boundaries (limiting the ways of how the user can interact with the world). Boundaries are not explicitly authored, since they are implied in the authored content. Rather, for the author the key aspect to keep in mind is to find creative ways to justify the existing boundaries to the players. For example, Façade [14] sets up a context (invitation for a dinner) that justifies the spatial boundary (all action takes place in one room) set by the authors.

Critical Mass for Emergence. Within well-defined boundaries, the authoring of content material is meant to “cover” the story landscape. As in any emergent system, a certain critical mass in terms of content is necessary for interesting narratives to emerge. This critical mass is not in absolute terms of quantity, but in relative terms of density, i.e., how well the authored content serves to create different paths through the narrative landscape. It is hard to find out whether the critical mass has been reached other than by playtesting and authoring. It is however important when designing content that the author keeps the density aspect in mind and does not view achieving the “critical mass” as a purely quantitative aspect. If a particular piece of added content adds new possibilities but also widens the boundaries of the story landscape, the density can go down rather than up; this is detrimental to the achievement of the critical mass.

Dead Ends. We consider dead ends to be states in the story landscape where the emerging narrative ends, not because the story has reached “the end”, but because there is a lack of content (e.g., an authored character goal with no actions to attain it, or characters moving in different directions so necessary interactions are no longer possible). We suggest that authoring for EN is a continuing process involving finding dead ends and resolving them by authoring new content for that situation. An open issue for this process is the question how to *detect* such dead ends. A promising approach might lie in automated tools that run the EN many times and try to construct a representation of the story landscape (similar to functionality in the Storytron engine [5]). Assuming the existence of such a tool, an author still needs to be aware that there is no direct relation between density and the amount of dead ends. Having no dead ends does not necessarily imply a high density or having reached the critical mass.

4.2 EN Simulation: Modelling a Dramatic Abstraction

The EN approach to storytelling is essentially one of *simulation*. We use Frasca’s somewhat broad definition of simulation being the “act of modelling a system A by a less complex system B, which retains some of A’s original behaviour.” [15]. Here, we consider system A to be the author’s envisioned dramatic world and system B to be the EN system.

A consequence of this notion of simulation is that the development of character models used in an EN architecture must be seen as a process of authoring, rather than as ongoing cognitive modelling research, and although the two might inform each other, research in this direction will not lead to “the” EN character model. In this sense, using simulation for storytelling requires a deep reflection by the author on the story world. One needs to not only think of what the envisioned characters are supposed to do, but also make explicit *why* they do what they do. This may seem like a laborious and excessive endeavour from a design point of view, but it pays off in terms of what is mediated. Interactors can engage with these simulated characters and discover, through playing, the story world and the rules governing it that the author wants to mediate.

The word ‘simulation’ brings along many connotations from its more prevalent and established use in research. By making these connotations explicit, we can explain certain pitfalls that must be avoided.

First of all, simulation in many cases means building models of (an aspect of) reality. In EN, we model a *dramatic abstraction* of reality. This means the character models in EN are based on how the author envisions characters to behave, rather than on how people behave in reality. This is not to discredit the use of established cognitive agent models to model EN characters, but rather to say that trying to make characters adhere to realistic behaviour is an *authorial choice*. Modelling after realism can be a valuable approach for writing the simulation (just like traditional authors will do their research to make their fiction resemble aspects of reality), and there might be good reason for doing so (e.g., to increase the suspension of disbelief, or to simulate realistic consequences for educational purposes), but we found that it can be very tempting to delve into

cognitive modelling and lose awareness of the ultimate purpose (which is the envisioned experience).

For instance, FAtiMA (the agent architecture used in FearNot!) models how characters respond emotionally to the prospect of their goals succeeding or failing, how characters select which goals to pursue, and how they decide how to act in a way that seems motivated by their goals. Although FAtiMA and its ideas can be reused, one might make quite different modelling choices for the simulation of virtual Teletubbies [1] or even alien worlds [16].

Second, modelling implies complexity reduction. It is tempting to do this by generalization: modelling general actions and general emotions in an attempt to cover a broad range of possible stories with a small set of content. This happened for instance in the beginning of the Virtual Storyteller project. The resulting stories are stripped of any of the particularities that make stories so engaging. In our view, there are better ways to achieve complexity reduction: by defining clearer and smaller boundaries for the story landscape, and by only making abstractions where they suit the envisioned interactions.

4.3 Issues with Story Space Authoring

In creating content, the author might naturally try to run the imagined simulation mentally, which means that narratives do not emerge as a surprise from the system itself but also materialize within the author's mind using their own understanding of the workings of the system. While this is not necessarily a problem, our experience with FearNot! has shown that this limits the emergent quality of the simulation. Once the author tries imagining the outcome of the narrative as the sole inspiration for modelling, there is a risk of "thinking in terms of plot" and the temptation to take a top-down view on things. The result is that the author starts following a narrow story path and creates just those elements that will produce this path, which has a negative effect on the density of the story landscape. There is no easy solution to this problem. After all, clinging to a plot helps the authors in covering the story landscape more consistently and avoiding dead ends, however with the result of narrowing it down.

In one strand of our work, we are currently exploring a possible approach to this authoring issue by means of massively collaborative authoring [17]. The hypothesis of this approach is that if the content for the simulation is provided by a group of authors, an individual author cannot predict or control the outcome of the simulation any more, creating a collaborative "letting go" attitude towards authoring. Local interactions between the content provided by different authors might create surprising emergent situations that an individual author could not produce, and author contributions can work mutually inspiring. After all, if we are assuming many different users exploring the dramatic world offered by an emergent narrative in many different creative ways it seems reasonable to assume that the creativity to create this world also needs to come from many authors. And besides the qualitative and creative aspect, having multiple authors also helps with the quantitative aspect (i.e., the aforementioned critical mass).

5 Conclusions

In this paper, we focused our reflection on identifying the actual processes that allow for an emergent narrative to take place. Along with the internal storification process of the interactor, we looked at the how an unfolding story affects the range of choices offered to a character. In doing so we identified that the role played by an interactor is more important than previously thought and that the responsibilities of a compelling and meaningful emergent performance do not only depend on the system but also on the actions of the interactor. In addition to looking at the role of a participant in an EN drama, we also started to formalize the role of the author and reflected on how author and interactor roles actually fit with the notion of narrative agency, by exploring the role an EN can take as a medium for storytelling. We argued that if the interactor is willing to play along with the EN system, an EN can mediate story-like experiences that are more personal and more nuanced than what can be conveyed through traditional storytelling. Finally, we considered EN authoring with respect to both interactor and author and identified the main issues that an author faces when creating an EN.

In this paper we aimed to position our research so as to set the basis for investigating EN further, especially in terms of real-time management, narrative development and authoring. We aim in time to paint a clearer picture of the concept that could be directly translated into applications that are entertaining and/or educational in offering interactors a story-like experience. While there are still many questions unanswered, we aimed with this paper to establish research directions that would lead to the advancement of knowledge in the IS domain and the more specific area of EN.

Acknowledgements

This research has been supported by the GATE project, funded by the Netherlands Organization for Scientific Research (NWO) and the Netherlands ICT Research and Innovation Authority (ICT Regie), and by the eCIRCUS (Contract no. IST-4-027656-STP) project carried out with the provision of the European Community in the Framework VI Programme. The authors are solely responsible for the content of this publication. It does not represent the opinion of the European Community, which is not responsible for any use that might be made of data appearing therein.

References

1. Aylett, R.: Narrative in virtual environments - towards emergent narrative. In: Proceedings of the AAAI Fall Symposium on Narrative Intelligence, pp. 83-86 (1999)
2. Aylett, R., Louchart, S., Dias, J., Paiva, A., Vala, M., Woods, S., Hall, L.: Unscripted narrative for affectively driven characters. *IEEE Computer Graphics and Applications* 26(4), 42-52 (2006)

3. Louchart, S., Aylett, R.: Emergent narrative, requirements and high-level architecture. In: Proceedings of the 3rd Hellenic Conference on Artificial Intelligence (SETN 2004), pp. 298 – 308 (2004)
4. Swartjes, I., Vromen, J.: Emergent story generation: Lessons from improvisational theater. In: Intelligent Narrative Technologies: Papers from the AAAI Fall Symposium (2007)
5. Crawford, C.: Chris Crawford on Interactive Storytelling. New Riders Games (2004)
6. Laurel, B.: Computers as Theatre. Addison-Wesley Publishing Company, Inc., Reading (1991)
7. Aylett, R., Louchart, S.: Being there: Participants and spectators in interactive narrative. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 117–128. Springer, Heidelberg (2007)
8. Kelso, M.T., Weyhrauch, P., Bates, J.: Dramatic presence. PRESENCE: The Journal of Teleoperators and Virtual Environments 2(1), 1–15 (1993)
9. Swartjes, I., Theune, M.: The virtual storyteller: Story generation by simulation. In: Proceedings of the 20th Belgian-Netherlands Conference on Artificial Intelligence (BNAIC) (2008)
10. Louchart, S., Aylett, R.: From synthetic characters to virtual actors. In: Proceedings Third Artificial Intelligence and Interactive Digital Entertainment, AIIDE 2007 (2007)
11. Swartjes, I., Kruizinga, E., Theune, M., Heylen, D.: Emergent narrative and late commitment. In: Prendinger, H., Lester, J.C., Ishizuka, M. (eds.) IVA 2008. LNCS, vol. 5208, pp. 543–544. Springer, Heidelberg (2008)
12. Murray, J.H.: Hamlet on the Holodeck: the Future of Narrative in Cyberspace. MIT Press, Cambridge (1998)
13. Young, R.M.: The cooperative contract in interactive entertainment. In: Dautenhahn, K., Bond, A.H., Canamero, L., Edmonds, B. (eds.) Socially Intelligent Agents: Creating Relationships with Computers and Robots, pp. 229–232. Springer, Heidelberg (2002)
14. Mateas, M., Stern, A.: Façade: An experiment in building a fully-realized interactive drama. Technical report, Literature, Communication and Culture and College of Computing, Georgia Tech. (2003)
15. Frasca, G.: Simulation 101: Simulation versus representation (2001) (Retrieved 13, March 2008), <http://www.ludology.org/articles/sim1/simulation101.html>
16. Kriegel, M., Lim, M.Y., Nazir, A., Aylett, R., Cawsey, A., Enz, S., Rizzo, P., Hall, L.: ORIENT: An inter-cultural role-play game. In: Proceedings of Narrative in Interactive Learning Environments (NILE) 2008 conference (2008)
17. Kriegel, M., Aylett, R.: Emergent narrative as a novel framework for massively collaborative authoring. In: Prendinger, H., Lester, J.C., Ishizuka, M. (eds.) IVA 2008. LNCS, vol. 5208, pp. 73–80. Springer, Heidelberg (2008)

From Debugging to Authoring: Adapting Productivity Tools to Narrative Content Description

David Pizzi and Marc Cavazza

School of Computing, University of Teesside
TS13BA Middlesbrough, United Kingdom
{D.Pizzi,M.O.Cavazza}@tees.ac.uk

Abstract. Recent progress in Interactive Storytelling has been mostly based on the development of proof-of-concept prototypes, whilst the actual production process for interactive narratives largely remains to be invented. Central to this effort is the concept of authoring, which should determine the relationships between generative technologies underlying Interactive Storytelling engines, and the actual description of narrative content. In this paper, we report the development of an authoring technology on top of a fully-implemented Interactive Storytelling system. Although this system originated as a debugging tool for a Planning system, its interactivity as well as the high-level nature of the formalism it manipulates makes it a candidate to support collaboration between authors and technologists.

Keywords: Authoring Tools, Content Creation, Character-based Storytelling, Planning.

1 Introduction

Whilst the idea of Interactive Storytelling (IS) can be traced back to the seventies, recent progress over the past 10 years can be attributed to the adoption of generative techniques based on Artificial Intelligence (AI) formalisms. One of the principal manifestations of progress has been the multiplication of proof-of-concept IS prototypes [38][21][20][26][25], which have often required considerable effort not just in terms of implementation but also content description. The development of these prototypes has often taken place in its entirety within computing research laboratories: in the absence of a declared interest by potential users (i.e. authors), example narratives have been developed by researchers themselves, often borrowing simple plots from tales, classics or popular culture (e.g. Façade and the Edward Albee's *Who's Afraid of Virginia Woolf* [21]). If we believe IS has reached a certain maturity, it can only refer to the viability of underlying computational techniques. The issue of content creation is lagging behind the early technical developments and their proof-of-concept prototypes. Its cornerstone is likely to be the relation between content creation and the generative formalisms making interactivity possible. This relation is generally described through the concept of *authoring*, which accounts for multiple aspects of the relation between content and technology. These range from the

creative and aesthetic aspects that would be specific to IS, to the more elementary ability of representing narrative content using the AI formalisms of IS prototypes. The growing interest in authoring in the IS community [32] finds its roots largely in the imbalance between the number of technical proof-of-concept prototypes and the number and scale of actual Interactive Narratives.

In this paper, we introduce an approach to authoring which evolved from the development of proof-of-concept prototypes and the need for internal productivity tools. We suggest how these tools can be adapted to a possible production process for IS, which is inspired from the development of computer games and more specifically the relationships between designers and programmers, yet operates at a higher level of abstraction, due to the ability of IS systems to directly formalise narrative actions. We first survey current approaches to IS authoring; after introducing our IS technology, we describe an associated authoring method and its embedding in an authoring software.

2 Related Work

The authoring process¹ remains a bottleneck limiting the further development of IS and the production of interactive narratives. As a result, the IS community has dedicated significant efforts to the identification of requirements for authoring systems and associated methods. Essentially, IS systems require the author to create/generate a large amount of appropriate content [21][20] to serve as a basis for narrative generation. Recently, Medler and Magerko [23] have proposed a list of essential requirements for an Interactive Drama Authoring Tool; they extracted six properties: *Generality* (it should be independent of both game environment and story content), *Debugging* (check for redundancy, dead-ends, poor perplexity in certain decision point, consistency, etc.), *Usability* (ease of learning, efficient to use), *Environment representation*, *Pacing and Timing* and, finally, *Scope*. Conversely, McNaughton et al. [24] identified four major problems being caused by the generative property of their system ScriptEase: *Generality*, *Performance*, *Coverage* and *Evolution*, which could in turn feed back into the definition of an authoring mechanism remediating these.

Previous work has proposed a classification of authoring tools: Szilas [34] associates authoring tools with the paradigm used by the IS engine itself. He classified existing tools amongst three distinct categories: “*character based*”, “*plot based*” and finally the last approach which consists in focusing on narrative properties rather than on a course of events or actions.

We propose a classification illustrating the duality between interfaces based on formalism and those based on story representation. The former foster generative aspects but lack content creation representation, the latter offer an intuitive visualisation of the story space and action sequences but force the author to explicitly express all the story alternatives limiting the plot variations.

¹ In this section, we deliberately exclude from our study authoring tools focused on the creation of visual elements to concentrate exclusively on the creation of narrative content.

Figure 1 proposes a classification of previous authoring systems along two main dimensions: *Visibility* and *Generativity*. Systems such as INSCAPE [39][7] and U-Create [28] use story-graph representation which offer a high visibility on the story plot but little generativity. On the other hand, system like ScriptEase [24] use scripted sequences of actions triggered by events which possesses substantial generative power but lacks visibility. The Scenejo authoring tool [33][37] differs sensibly from other systems as it is rather used for the generation of dialogues constitutive of a narrative in a way similar to Façade [21]. To the exception of most recent systems such as Scribe [23], Wide-Ruled [30] or even the authoring environment of the Bowman planning system [35], we can notice that most authoring systems tend to favour one property over the other.

Ideally, an authoring tool should both give access to the generative power of its formalism and allow the representation of the created narrative content. In our previous work [4], we investigated the fundamental relations between planning formalisms and authoring in IS. Despite its high power of generation, HSP [2] is often based on STRIPS formalism [9] which clearly lacks visibility for long-term dependencies between the actions carried out during the plan unfolding. To address this challenge, we have developed, as part of later experimentations, tools that will provide the representation of all the possible unfolding plans ensuring thus a greater visibility on the whole story space.

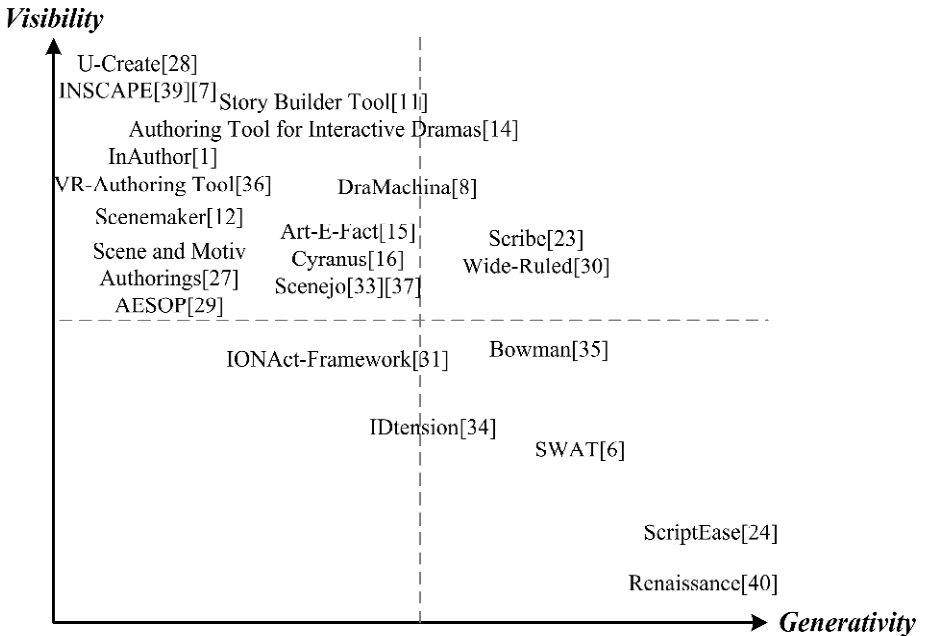


Fig. 1. Classification of existing authoring tools according to their visibility and generativity

3 Overview of Authoring Tools

We have decided to use a classic XIXth century French novel, *Madame Bovary* [10] by Gustave Flaubert, as the background for our interactive storytelling experiments. As described by Flaubert himself in his preparatory work to his novel *Madame Bovary* [19], characters' feelings can be considered as the missing link between the cognitive and a narrative perspective on the story's characters [5] and thus can be used to drive the development of the story plot. Combined with ontology of characters' feelings, our latest work on character-based affective IS [25] aims at reconciling both philosophies and uses a real-time Heuristic Search Planner (HSP) [2] [17] to generate characters' actions consistent with their psychology and allowing anytime user interactions.

Our authoring environment (Figure 2) has been developed subsequently to our proof-of-concept prototype of emotional planning for IS. Its rationale was to support the authoring of a complex planning domain, by checking its completeness and its consistency. However, since this authoring tool was an interface to the narrative formalism itself, and that the narrative formalism determined entirely the interactive narrative it became a candidate for a more generic approach to authoring. The integration of new modules could support the collaboration between authors and developers in designing an interactive narrative.

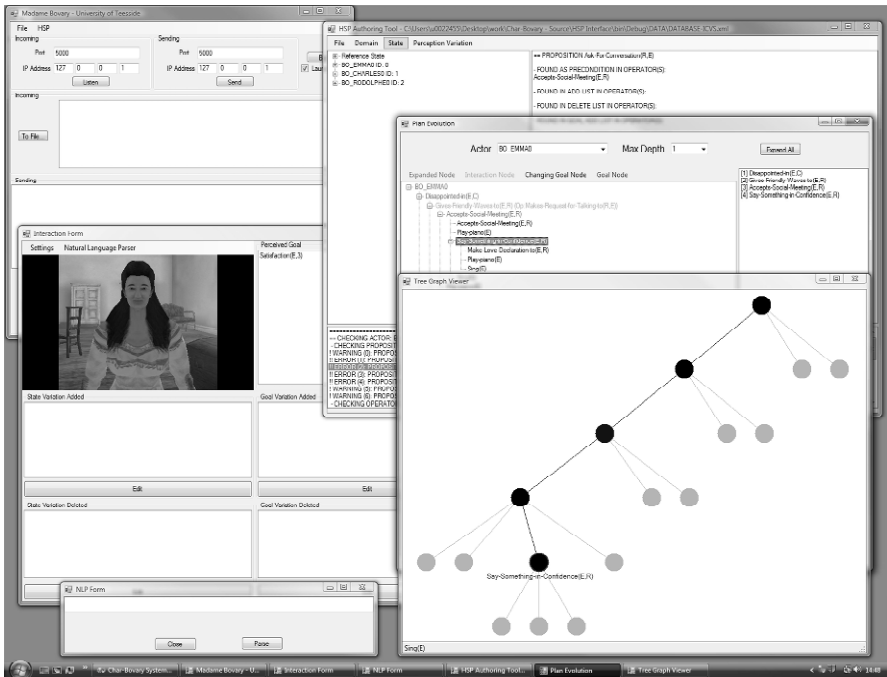


Fig. 2. Overview of our IS environment (run-time interface, interaction and authoring tools)

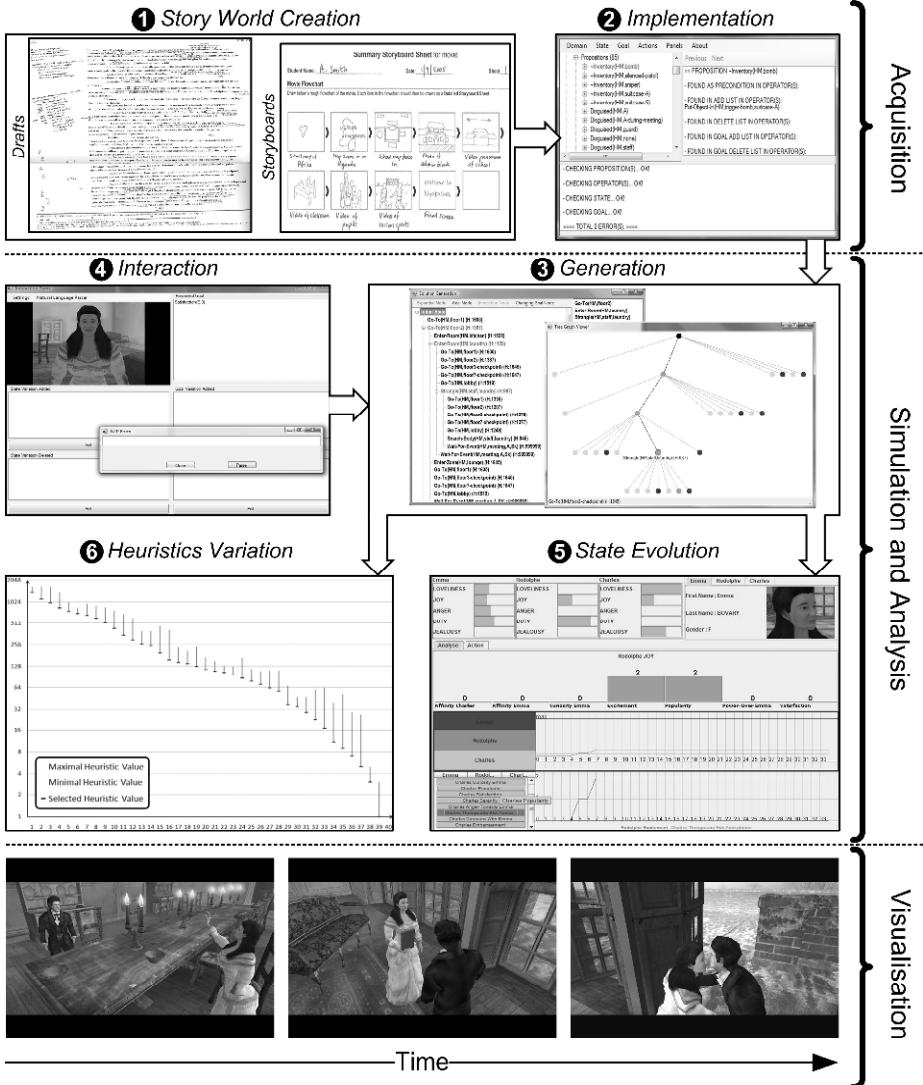


Fig. 3. Overview of the Authoring process and its three main stages (Knowledge Acquisition, Simulation and Analysis and Story Visualisation)

As illustrated by Figure 3, the authoring approach supported by our system follows three main stages: *knowledge acquisition*, *simulation and analysis*, and finally *story visualisation*.

Knowledge Acquisition. The creation of a story world (Fig 3-1) is the initial stage where drafts of story elements are created by the author. They describe diverse story elements (e.g. characters' psychology, representative scenes or environment description). The next step (Fig 3-2) corresponds to the elicitation of all knowledge

required to describe the story world, such as the various states (i.e. initial and goal states) and the various actions described through their validity conditions and their consequences. In terms of Planning, this corresponds to *domain implementation* where each part of the planning domain is created (i.e. propositions, operators, states and goal). The domain description also includes a formalisation of the initial state and the goal state, which correspond to the scene's or characters' objectives.

Simulation and Analysis. As an alternative to the offline generation of a complete narrative (formally a solution plan), an interactive mode allows a step-by-step generation of a solution including the visualisation of all possible outcomes (Fig 3-3). Starting from the initial state, the user can expand the plan at each step using a tree representation until the goal state is reached. After each action is selected by the user, the system automatically offers a list of possible subsequent actions. For instance, the system will only offer the solution of Emma accepting an invitation once Rodolphe would have proposed it. This simulation has both a formal component (access to the planning domain, inspection of operators and world states) and a visual component (a tree structure providing a natural visualisation of actions and their consequences). This dual visualisation is meant to support collaboration and explanation between system developers and content creators.

In addition, authors can interact with the solution generation process at any time (Fig 3-4). This module includes also a *dynamic environment simulation* feature. It allows reproducing changes in the world not triggered by the planning system that will normally occur within the story, without having to simulate this in the complete 3D environment.

Then, several analysis tools (Fig. 3-5 and 3-6) can allow the validation of the generated narrative content. For instance, the evolution of the world state through the development of the story plan is an effective way of ensuring its consistency with respect to the characters' psychology by allowing the analysis of how emotions intensities vary along the plan evolution.

Story Visualisation. Finally, when the result has been validated, the last stage is to visualise the final result using the run-time engine.

We can observe that the early step of this production process is a specific case of *knowledge engineering* applied to planning formalisms (i.e. by integrating knowledge into computer systems in order to solve complex problems normally requiring a high level of human expertise).

In the following sections, we develop the above mentioned stages, and illustrate them using an example.

4 Domain Implementation

The first step is the elicitation of all the planning knowledge required and starts by providing a complete propositional representation of the world (e.g. characters' beliefs such as $\text{Affinity}(X, Y, \text{intensity}) \mid \text{intensity} \in \{1, 2, 3, \dots, n\}$, contextual or interactions flags (e.g. $\text{Is-in-Intimacy}(X, Y)$), etc.).

The initial and goal states are then represented by conjuncts of propositions. The planning operators are represented using a STRIPS-like formalism (i.e. a set of propositions as preconditions and effects) [9] and correspond to actions that can be performed by the characters (e.g. *Is-Kissed-by*(E,R), *Physical-Contacts*(E,R) or *Feels-Guilty-of-Estranging*(E,C)).

Usually, the grounded sequence of elementary actions that will be played by the game engine on the operator selection is also determined at this stage.

The particular choice made on the granularity of the domain representation is not without any impact on the types of solutions that will be generated later. Common pitfalls consist in representing the domain at a level of abstraction too high or conversely using too specific actions, which will limit the generative power of the system. Indeed, a too low level of abstraction would only produce minimal story variations and not real compelling plot twists.

The description of the planning domain is a manual process which becomes error-prone when its size (and the corresponding number of operators) increases. It is one well-known problem to maintain consistency between the predicates used by the various operators, which can become challenging when describing operators manually. Inconsistencies in predicates' labels could also be responsible for errors in the content of operators with other detrimental side-effects. There is thus a need to check consistency of preconditions and effects every time changes to the planning domain are introduced as part of the knowledge elicitation process.

5 Plan Generation

When solutions rely on a sophisticated plan, the various causal dependencies as generated by HSP planning may be difficult to recognise. Therefore, we wanted to explore whether the set of possible plans could be visually represented in order to control the unfolding of the generated content. Moreover, the combinatorial aspect of content generation can quickly overflow the amount of possible paths that can be exploited if done using a brute force approach. For instance, if we consider the first reference milestone set by Mateas [22]: an interactive story is compelling if it presents a 10 minute length with one beat per minute. With an average of one beat every 4 actions (~15 seconds duration per action) it would thus represent a 40 steps long story. Now, with an average perplexity of 5 operators per step, a 40 steps long plan would simply present a total of 5^{40} (~ 10^{28}) possible solutions. For that reason, our system proposes a step-by-step plan simulation, in which the user can manually expand nodes and from there visualise the results using a tree-like hierarchy (see Figure 4).

Every operator in the planning domain is tested for applicability. Whenever an operator preconditions are satisfied, it is applied to create a new state that could further be extended. For efficiency purposes, and in order to avoid redundancy, we only develop new states that diverge significantly from their parent state. The whole tree can then be automatically scanned and all the possible solutions can be listed and visualised.

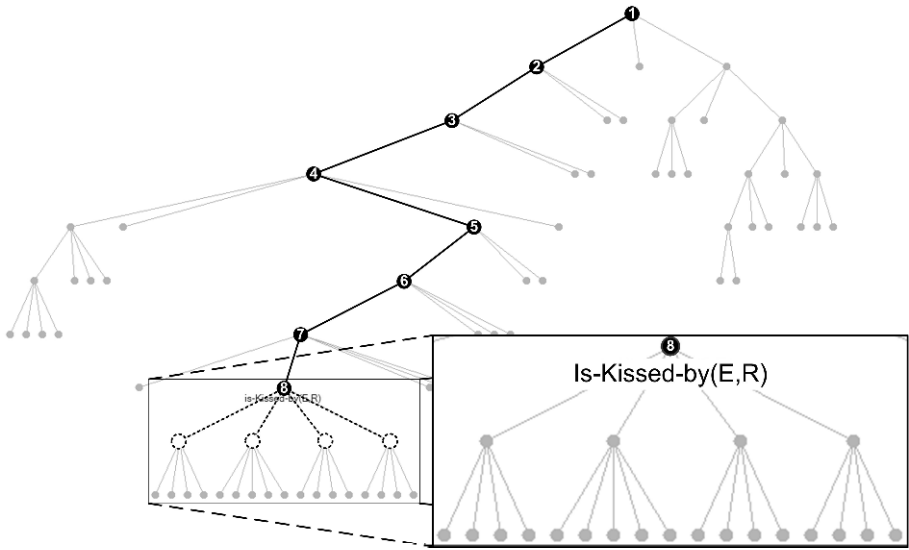


Fig. 4. Tree-based visualisation of the unfolding plan for the narrative

6 Simulating User Interaction during Authoring

We can specify interactions at any time in the form of applied operators or more simply by directly editing the propositions present in the world state. The tool generates from there the new unfolding plans saving a significant time during the domain creation.

Figure 5 shows the impact of an interaction upon the following action selection. In the original plan and once the action *Is-Kissed-by*(E,R) has been executed by the game engine, Emma will have to choose her next action amongst four different operators. According to the best heuristic value [18], the HSP planner will select the operator *Physical-Contacts*(E,R) has it presents a minimum heuristic of 19. Now, it is obvious that one user interaction could induce a variation in the following available actions and thus on the subsequent plan generation. Using the tool allowing the input of user utterance, we enter the sentence “*you shall not be one of these frivolous women*” which will have the effect of adding the proposition *Embarrassment*(E,3) within Emma’s beliefs (see [25] for more details). As a consequence, we can observe that following the action *Is-Kissed-by*(E,R) five different operators can now be applied. We can also see that the best subsequent heuristic value (6) belonging now to an alternative operator *Feels-Guilty-of-Estranging*(E,C), the HSP planner will produce an alternative plan that will lead Emma to return to her husband forgetting about her new love affair.

In a real game engine, the environment is often dynamic (e.g. Non-Player Character (NPC) have autonomous behaviour). For instance, a NPC can keep walking from one room to another on cycles. Consequently, spatio-temporal variations have also to be simulated within the solution generation process. We have introduced the

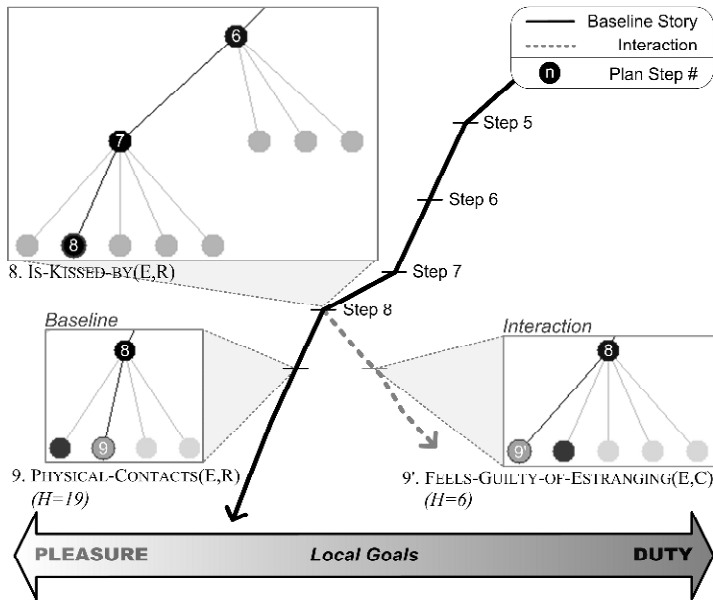


Fig. 5. Representation of the impact of user interaction upon the unfolding plan

possibility to modify the world state at any stage in order to make the appropriate state variations. The plan generation is synchronised using external actions with plan interruptions reproducing characters movements (e.g. Update-Position) or world events (e.g. Character-Level-Arrival).

7 Analysis and Validation

In line with our commitment to the exploration of aesthetic/dramatic phenomena in the narrative, we are looking for a representation that would describe the overall story progression. The direct analysis of actions' sequence seems to be a simple method of appreciating consistency and can thus constitute a preliminary validation of the generated narrative.

However, if we seek inspiration from narratology, Bremond [3] has proposed to evaluate the characters' situation in terms of improvement and deterioration, including multiple causal descriptions of the reciprocal characters' influence on the process. Dramatic aspects can then occur from such influences and improve the overall plot quality. Reasoning on causal dependencies in plans seems better suited to task models [26] or plan-space search, in which such dependencies can be explicitly formalised (see e.g. the examples discussed in [13]). For state-space search, as implemented when using HSP, the contribution of specific actions to an agent's desires is best measured through the heuristic function. For example, analysing the evolution of the ranges of heuristic function variations can inform about the diversity and quality of story alternatives. The evolution of the world state composition can also inform authors about how the emotional dimensions are exploited during the plan, which can

help to ensure the believability and consistency of the characters according to their psychology.

Therefore, the generation of narrative content should not remain the only areas that are assisted by authoring tools. Our initial productivity tools have been refined to support both the complete analysis of generated sequences (Section 5) and the impact of anytime user interactions upon the unfolding story (Section 6). The formalisation of narrative actions into planning operators (STRIPS or PDDL) undoubtedly requires some knowledge of Planning in both description and appraisal of the domain. This limits the use of such tools to either programmers or authors motivated enough to acquire knowledge engineering skills. However, if the ability of the IS community to produce validation method is critical for its maturity, we believe that the experience gained at expert level could constitute a first case study and a starting point to the development of authoring methodologies for designers and writers.

8 Conclusion

We have presented our development of an authoring technology on top of our fully-implemented IS system based on Flaubert's novel *Madame Bovary*. Although this system originated as the debugging tool for our Planning system, its interactivity as well as the high-level nature of the formalism it manipulates makes it a prominent candidate for the support of collaboration between authors and technologists.

One essential question that we need to ask ourselves as developers is whether an authoring tool could exist as fully independent from these underlying IS technologies. The idea of devising perfectly generic narrative formalisms may remain utopian as it is very difficult to extract them from the procedural concepts of narrative content generation. One of the demonstrations of this has been the proliferation of different authoring tools which have been developed over the past few years within the IS community.

Acknowledgments. Jean-Luc Lugin is thanked for the development of the visualisation environment.

References

1. Barrenho, F., Romao, T., Martins, T., Correia, N.: Authoring environment: Interfaces for creating spatial stories and gaming activities. In: Proceedings of the 2006 ACM SIGCHI international conference on advances in computer entertainment technology (2006)
2. Bonet, B., Geffner, H.: Planning as heuristic search: new results. In: Biundo, S., Fox, M. (eds.) ECP 1999. LNCS, vol. 1809, pp. 360–372. Springer, Heidelberg (2000)
3. Bremond, C.: *Logique du Récit*. Editions du Seuil, Paris (1973)
4. Charles, F., Lozano, M., Mead, S.J., Bisquerra, A.F., Cavazza, M.: Planning formalisms and authoring in interactive storytelling. In: Göbel, S., et al. (eds.) Proc. TIDSE 2003, pp. 216–225. Fraunhofer IRB Verlag (2003)
5. Christian, D., Young, R.M.: Comparing Cognitive and Computational Models of Narrative Structure. In: The Proceedings of the National Conference of the American Association for Artificial Intelligence (2004)

6. Crawford, C.: Storytron Interactive Storytelling, Project (last accessed July 1, 2008), <http://www.storytron.com/>
7. Dade-Robertson, M.: Visual Scenario Representation in the Context of a Tool for Interactive Storytelling. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 3–12. Springer, Heidelberg (2007)
8. Donikian, S., Portugal, J.: Writing interactive fiction scenarii with draMachina. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 101–112. Springer, Heidelberg (2004)
9. Fikes, R., Nilsson, N.: STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2, 189–208 (1971)
10. Flaubert, G.: Madame Bovary. In: revue de Paris, L. (ed.), France (1856) (in French)
11. Fuertes, J., Gonzalez, A., Mariscal, G., Ruiz, C.: Developing Virtual Storytellers for the Virtual Alhambra. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 63–74. Springer, Heidelberg (2007)
12. Gebhard, P., Kipp, M., Klesen, M., Rist, T.: Authoring Scenes for Adaptive, Interactive Performances. In: AAMAS-2003. Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems, pp. 725–732 (2003)
13. Gratch, J.: Why you should buy an emotional planner. In: Proceedings of the Autonomous Agents 1999 Workshop on Emotion-based Agent Architectures (EBAA 1999) (1999)
14. Hartmann, K., Hartmann, S., Feustel, M.: Motif Definition and Classification to Structure Non-linear Plots and to Control the Narrative Flow in Interactive Dramas. In: Subsol, G. (ed.) ICVS-VirtStory 2005, vol. 3805, pp. 158–167. Springer, Heidelberg (2005)
15. Iurgel, I.: From Another Point of View: Art-E-Fact. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 26–35. Springer, Heidelberg (2004)
16. Iurgel, I.: Cyranus – An Authoring Tool for Interactive Edutainment Applications. In: Pan, Z., Aylett, R.S., Diener, H., Jin, X., Göbel, S., Li, L. (eds.) Edutainment 2006. LNCS, vol. 3942, pp. 577–580. Springer, Heidelberg (2006)
17. Korf, R.E.: Real-time heuristic search. *Artificial Intelligence* 42(2-3), 189–211 (1990)
18. Liu, Y., Koenig, S., Furcy, D.: Speeding Up the Calculation of Heuristics for Heuristic Search-Based Planning. In: The Proceedings of the Eighteenth National Conference on Artificial Intelligence, pp. 484–491 (2002)
19. Leclerc, Y. (Eds.): Plans et Scenarios de Madame Bovary, CNRS Edn., France (1995) (in French)
20. Magerko, B.: Building an Interactive Drama Architecture. In: Göbel, S., et al. (eds.) Proc. TIDSE 2003, pp. 226–237. Fraunhofer IRB Verlag (2003)
21. Mateas, M., Stern, A.: Façade: An Experiment in Building a Fully-Realized Interactive Drama. In: Game Developer's Conference, San Francisco, CA (2003)
22. Mateas, M., Stern, A.: A Behavior Language for Story-Based Believable Agents. *IEEE Intelligent Systems* 17(4), 39–47 (2002)
23. Medler, B., Magerko, B.: Scribe: A Tool for Authoring Event Driven Interactive Drama. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 139–150. Springer, Heidelberg (2006)
24. McNaughton, M., Cutumisu, M., Szafron, D., Schaeffer, J., Redford, J., Parker, D.: ScriptEase: Generative Design Patterns for Computer Role-Playing Games. In: ASE 2004 (2004)
25. Pizzi, D., Charles, F., Lugin, J.-L., Cavazza, M.: Interactive Storytelling with Literary Feelings. In: Paiva, A.C.R., Prada, R., Picard, R.W. (eds.) ACII 2007. LNCS, vol. 4738, pp. 630–641. Springer, Heidelberg (2007)

26. Riedl, M.O., Stern, A.: Believable Agents and Intelligent Story Adaptation for Interactive Storytelling. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 1–12. Springer, Heidelberg (2006)
27. Rober, N., Huber, C., Hartmann, K., Feustel, M., Masuch, M.: Interactive Audiobooks: Combining Narratives with Game Elements. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 358–369. Springer, Heidelberg (2006)
28. Sauer, S., Osswald, K., Wielemans, X., Stifter, M.: U-Create: Creative Authoring Tools for Edutainment Applications. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 163–168. Springer, Heidelberg (2006)
29. Silverman, B., Johns, M., Weaver, R., Mosley, J.: Authoring Edutainment Stories for Online Players (AESOP): A Generator for Pedagogically Oriented Interactive Dramas. In: Balet, O., Subsol, G., Torguet, P. (eds.) ICVS 2003. LNCS, vol. 2897, pp. 65–73. Springer, Heidelberg (2003)
30. Skorupski, J., Jayapalan, L., Marquez, S., Mateas, M.: Wide Ruled: A Friendly Interface to Author-Goal Based Story Generation. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 26–37. Springer, Heidelberg (2007)
31. Sobral, D., Machado, I., Paiva, A.: Managing Authorship in Plot Conduction. In: Balet, O., Subsol, G., Torguet, P. (eds.) ICVS 2003. LNCS, vol. 2897, pp. 57–64. Springer, Heidelberg (2003)
32. Spierling, U., Iurgel, I.: Pre-Conference Demo Workshop “Little Red Cap”: The Authoring Process in Interactive Storytelling. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 193–194. Springer, Heidelberg (2006)
33. Spierling, U., Weiß, S., Müller, W.: Towards Accessible Authoring Tools for Interactive Storytelling. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 169–180. Springer, Heidelberg (2006)
34. Szilas, N.: IDtension: a narrative engine for Interactive Drama. In: Göbel, S., et al. (eds.) Proc. TIDSE 2003, pp. 187–203. Fraunhofer IRB Verlag (2003)
35. Thomas, J., Young, R.M.: Author in the Loop: Using Mixed-Initiative Planning to Improve Interactive Narrative. In: The ICAPS 2006 Workshop on AI Planning for Computer Games and Synthetic Characters (2006)
36. Wages, R., Grutzmacher, B., Conrad, S.: Learning from the Movie Industry: Adapting Production Processes for Storytelling in VR. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 119–125. Springer, Heidelberg (2004)
37. Weiß, S., Müller, W., Spierling, U., Steimle, F.: Scenejo – An Interactive Storytelling Platform. In: Virtual Storytelling, Using Virtual Reality Technologies for Storytelling, Third International Conference, Proceedings, Strasbourg, France, pp. 77–80 (2005)
38. Young, R.M., Riedl, M.O.: Towards an Architecture for Intelligent Control of Narrative in Interactive Virtual Worlds. In: Proceedings of the International Conference on Intelligent User Interfaces (IUI), Miami, FL, USA, pp. 310–312 (2003)
39. Zagalo, N., Göbel, S., Torres, A., Malkewitz, R.: INSCAPE: Emotion Expression and Experience in an Authoring Environment. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 219–230. Springer, Heidelberg (2006)
40. Zancanaro, M., Cappelletti, A., Signorini, C., Strapparava, C.: An Authoring Tool for Intelligent Educational Games. In: Balet, O., Subsol, G., Torguet, P. (eds.) ICVS 2001. LNCS, vol. 2197, pp. 61–68. Springer, Heidelberg (2001)

PRISM: A Framework for Authoring Interactive Narratives

Yun-Gyung Cheong, Yeo-Jin Kim, Wook-Hee Min,
Eok-Soo Shim, and Jin-Young Kim

Graphics&OS Group, Software Labs, SAIT, CTO, Samsung Electronics CO., LTD.
Giheung-Gu, Gyeonggi-Do, South Korea

{yuna.cheong, yeoj.kim, wookhee.min, esp.shim,
claire.j.kim}@Samsung.com

Abstract. The advances in computing technologies enable the computer users to create and share their own stories to the community at large. However, it is still regarded as complicated and laborious to author interactive narratives, where a story adapts as the user interacts with it. In authoring interactive narratives, two main approaches—branching graphs and AI planning—have been significantly used to augment interactivity into conventional linear narratives. Although each approach offers its own possibilities and limitations, few efforts have been made to blend these approaches. This paper describes a framework for authoring interactive narratives that employs an adapted branching narrative structure that also uses planning formalism to enable automated association between nodes. We expect that our work is valuable for non-expert users as well as AI researchers in interactive storytelling who need to create a large quantity of story contents for varied endings for a story.

Keywords: Interactive digital storytelling, authoring tool, interactive contents, story representation, branching narrative, story graph, AI planning.

1 Introduction

The advances in computing technologies enable the computer users to create and share their own stories to the community at large. This is evidenced by the significant success of animated movies and UCCs on the Internet, and the emergence of Machinima—the production technique of computer-generated animated movies using 3D game engine. Therefore, UCC providers such as YouTube supply video authoring tools for their users. However, it is still regarded as complicated and laborious to write interactive narratives, where a story adapts as the user interacts with it. Unlike conventional narrative systems in which one single story is told linearly without intervention from the audience, interactive narrative environments need to arrange multiple plots and unfold one of them as the audience makes a series of choices. For this, interactive story systems require an effective methodology to represent a collection of plots integrated with the audience's choices.

While a number of techniques for authoring interactive narrative generation have been presented by AI researchers and Game developers [1; 3; 4; 7; 9; 11; 12; 16; 21],

this paper investigates two main approaches to interactive story representation: branching narratives and plan structures. A branching narrative often takes the form of a directed graph containing nodes and arcs between the nodes; a node denotes a series of scripted scenes, and an arc denotes transition from one node to another.

The branching narratives have been favored in a number of practical interactive story systems due to its compelling advantages—intuitiveness and high degree of expressiveness [5; 16; 17; 18]. The interactive digital storytelling system *Scenejo* [16] provides an authoring tool that enables the author to construct a transition graph to represent plot progress where plot points are reached thorough dialogue exchanges between multiple agents. In the plot graph in *Scenejo*, a scene is represented as a knowledge base that describes the participating agents' utterances using a pattern matching technique. A transitional condition that advances the story to the next scene in *Scenejo* falls in three categories: elapsed time, a user utterance, or a state change. As exemplified in [16], guessing a right password may result in a scene change.

However, it has been also noted that an interactive story system employing a branching technique which needs frequent interaction with the user would face explosion of conditional branches. In addition, a branching narrative is typically built at the design time and thus its interactivity with the audience is limited by the story author's anticipation. As a result, the player interactions that are not considered at the design time often ignored and unhandled. Furthermore, the story viewers of identical choices would have same experiences with the system [14]. To address these issues, Iurgel [4] incorporates a rule-based technique into the authoring tool *Cyranus* which is embedded in an interactive story platform *Art-E-Fact*. In his system, some nodes exist outside of the story graph are labeled terminal nodes and can be activated by a rule-based engine without incoming transitions from other nodes in the graph.

In contrast, a narrative system employing planning formalisms takes as input a story which lacks in the user choice and adapts it as the user interacts with the system [2; 6; 8; 13; 20]. While the use of planning formalism in the interactive narrative system relieves the story author of integrating story contents with the user interaction, the system generated stories often give less immersive experience to the system users.

This paper discusses a framework for authoring interactive narratives that employs a hybrid approach to creating interactive stories. Our approach uses branching narrative structures which contain planning annotations (i.e., preconditions and effects) in order to enable automated association between nodes. This framework is a part of a project implementing a software package which targets non-expert users as well as experts in interactive storytelling. We plan to provide the package with rich 3D graphic models and other story elements such as sound effects and music so that even non-expert users can easily create highly engaging interactive stories. The framework utilizes the TVML (TV Program Making Language) technique [22] for story realization into 3D animation. TVML is a text-based language that can describe all the required elements (e.g., stage set, camera, light, characters, prop, voice, sound, etc.) for a 3D animation production. TVML scripts are realized as 3D animation on any TVML player equipped platforms (e.g., television sets, mobile devices, personal computers, etc.). Thus, we believe that this ubiquitous aspects of our framework will be also beneficial for conventional interactive narrative generation systems which typically need to deal with a large quantity of story contents to create stories with a great deal of variability [7; 15; 20].

2 The System Overview

The system provides a pair of authoring tools in the view of the authoring process of interactive narratives as involving two stages: content creation and interactivity creation. Figure 1 illustrates the system architecture involving one story writer, one interaction engineer, and one story participant, in which they could be a single individual. In the figure, the two components for interactive narrative together are called PRISM whose purpose is to take in the system user's behavior (i.e., input) and present diverse storylines as an analogy to a prism separating the light into different colors.

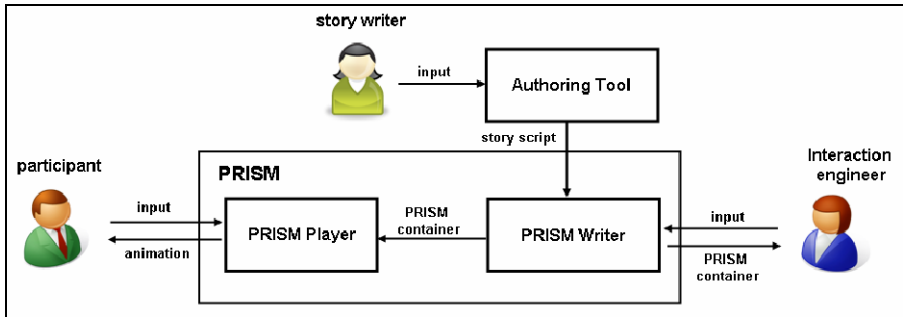


Fig. 1. The interactive story authoring architecture in which PRISM is situated

For content creation, the system provides a graphical user interface (as in Figure 2 in the next page) which employs the TVML technology as its rendering module. The interface allows the user to select rich 3D model items (e.g., stage sets, characters, props) and to control light, camera, music, actor behaviors, dialogues, narration, and captions. Because the content creation authoring tool is stand-alone, it is also available to the users for the exclusive purpose of creating non-interactive stories. The details of the content creating interface are beyond the scope of this paper.

On the completion of authoring story contents in the interface, the authoring tool generates a scene script file, which is an xml document that contains scenes, beats, and shots defined in PRISM as following. The script is then sent to the PRISM writer.

Scenes. A scene describes all the story elements that take place on the same stage.

Beats. A beat is an expressive unit in a story. A beat may include a series of shots that composes a meaningful event. For example, a beat describing a lecture on a specific topic may include shots of teacher explaining the topic followed by a shot of question and answer session.

Shots. A shot describes a set of story elements taken in one camera shot. Depending on the types of camera technique, a shot may span from one second to a few minutes. As an example of a long shot, a camera may track a young girl walking through a winding road in a province. In contrast, in a scene with two actors talking the camera may take short bust-shots of two actors who are facing each other to illustrate their conversation.

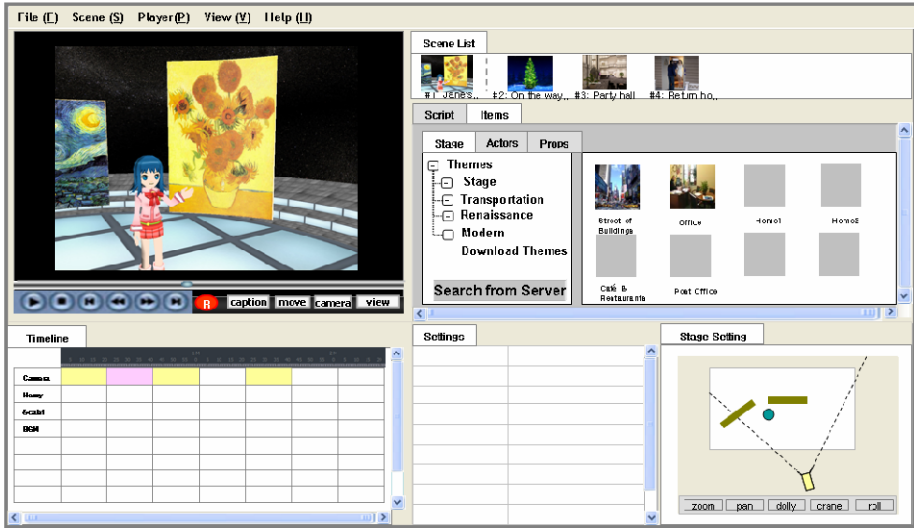


Fig. 2. Screenshot of the prototype User Interface window for content creation

3 The PRISM Writer

The PRISM writer enables the author to fabricate scene scripts into an interactive story utilizing the user interface shown in Figure 3 in the next page. The interface allows the author to construct an interactive story structure and encapsulates it in an XML format—a PRISM container—which is then passed to the PRISM player for presentation to the story participant.

3.1 The PRISM Writer Interface

An interactive narrative structure in PRISM is called a story map (see Figure 3) which represents a story space that the user may navigate and explore. In order to construct a story map, a concept of *situation* is introduced as its basic unit. A situation is a semantic story element that may include a number of beats in disparate scenes. Imagine, for example, a situation where an old man in an asylum recalls an excursion to a state fair with his family decades ago. This situation may contain two beats from two distinct scenes: a beat illustrating him recalling in an asylum scene and a beat illustrating a joyful excursion at a state fair scene. When a situation is created, the author can place the situation onto the story map area by the drag-and-drop style interface.

When given a scene script, the PRISM writer decodes the script to obtain scene information such as title, description, and the first screenshot of the scene. Next, the PRISM writer user interface lays out the scenes contained in the script on the Scene Area located on the right upper area of the window. The author then may click on one of the scenes to lay out the scene’s beats on the BeatSequence Area. Then, the author can compose a situation by dragging and dropping beats from the BeatSequence Area

onto the Situation Area. Although the beats and scenes are depicted as labeled rectangles in Figure 3, an implementation of the PRISM writer selects the first shot of each beat and scene as it delegates to show them on the BeatSequence and Scene areas for the author to easily recognize them.

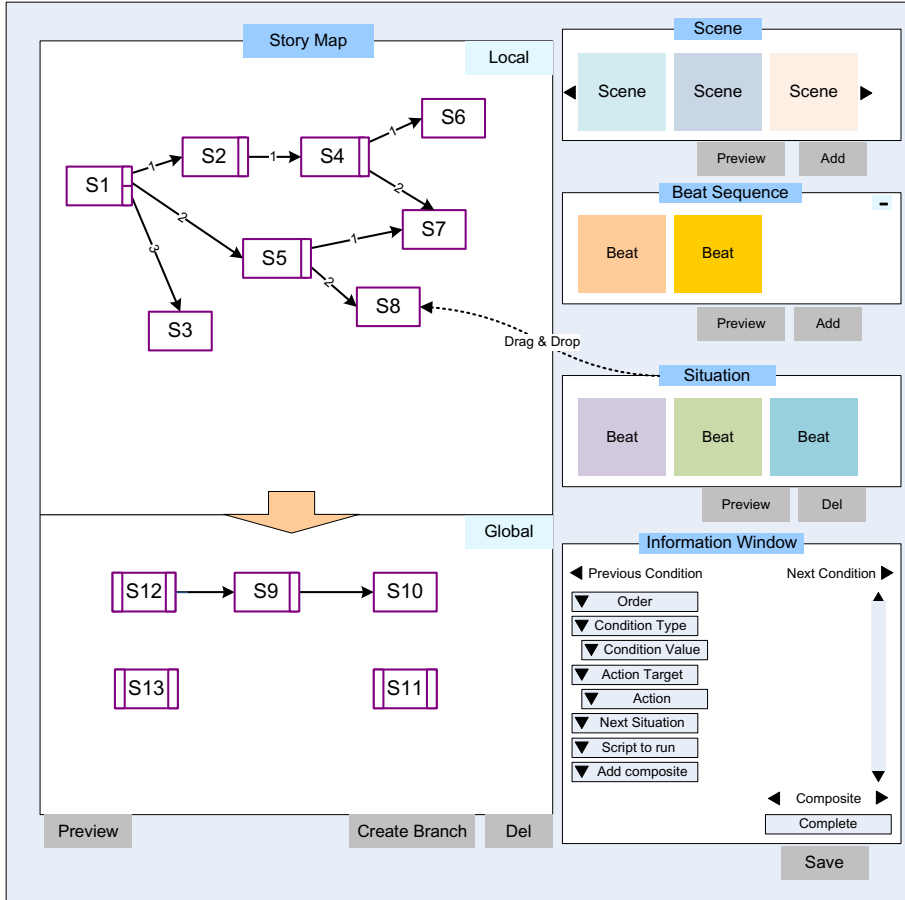


Fig. 3. The PRISM Writer Concept UI Design

The Global situation area on the left lower area of the PRISM writer interface (Figure 3) allows the author to specify a group of situations as global situations; global situations can occur at any point in the presentation of a story when their conditions are met. For instance, the author may produce supplementary situations to be provided when the user requests more information such as the background of the story. The author can easily convert global situations to local situations by dragging them into the local story map area.

3.2 The PRISM Container

The PRISM container is an xml document that describes a narrative structure which fuses branching graphs with planning formalism. As depicted in Figure 3, the story structure in PRISM basically takes a directed graph in which nodes represent situations connected by arcs representing the choice provided by the user or external systems (e.g., weather information, location data). While a conventional branching graph approach provides an intuitive and efficient way for creating interactive stories, it is also disadvantageous in that it restricts the user interaction to a set of actions anticipated by the story author. To compensate this drawback, PRISM also employs a planning approach. In the story graph in Figure 3, rectangles represents situations, where a bar (or a smaller rectangle) on the left denotes its precondition and a bar on the right denotes its effect. When the story graph contains internal nodes that contain no outgoing branches, our planner takes on the job to create links between nodes. As a result, in connecting situations where both approaches are available, PRISM gives higher priority to the branching approach regarding the author created branching as her direct intention. More details on the use of planning techniques in PRISM are discussed in the section 4.1.

In order to incorporate planning formalism into branching narratives, we devise a data structure named *condition* to be shared by both approaches. A condition is a logical sentence which describes a type of user interaction (e.g., text, speech, gesture, facial expression, mouse, brain signal, etc.), external information (e.g., real-time web data, GPS, etc.), and states of the virtual environment in which the story is situated (e.g., scores, characters' locations, the story participant's cognitive and emotional model, story progress, etc.). In the story graph, a condition can serve as a branching decision point that connects a single situation to another. Suppose, for instance, an educational application that displays a pair of fruit props, an apple and an orange, on the screen and waits for the user choice. The node that describes this example may include a condition that links her mouse-click event on the apple prop in a situation to another situation which explains planting apples in South Korea. In this example, the current situation contains a condition represented as a logical sentence (e.g., *User-Mouse-Clicked(apple)*) and embeds a branch that links to the explanation situation.

In the planning mode in PRISM, conditions can serve to represent preconditions and effects. Preconditions are logical statements that must be achieved for a situation to occur. Effects are logical statements denoting significant states that are altered in the virtual story world due to the situation execution. For this, a situation data structure holds places for preconditions and effects, which are optional. As an illustration, suppose that a penniless man who steals a bank with a shotgun. For this situation to take place, the action requires him holding a shotgun as its precondition (i.e., *Holding(man, shotgun)*) and describes him being rich as its effect (i.e., *Rich(man)*). To assist the user with constructing conditions, the interface provides the user with a predefined list of conditions and their combinations as well as a text editing box for writing customized conditions.

This hybrid representational approach makes it feasible for the user to reach disjoint sub-graphs as long as the graphs contain appropriate planning annotations—achievable preconditions from execution of other situations. If these two different approaches mutually inconsistent—a branch of a node points to *A node* while its

planning annotation leads to *B node*, the branching approach takes precedence over the planning, in the respect of the author’s intention. However, such conflicts would not happen in the current design since the planner is only active when narrative branches are not available. More details on the PRISM container can be found in [10].

3.3 The PRISM Writer Architecture

This subsection describes the PRISM writer architecture underlying the PRISM writer interface. The PRISM writer takes a scene script and generates a PRISM container corresponding to the story map on the PRISM writer interface. As illustrated in Figure 4, the PRISM architecture consists of four components: a story map creator, a post processor, an xml decoder, and an xml encoder. When the PRISM writer interface loads a scene script, the xml decoder analyzes the script and extracts scene and beat information. As the user makes modifications to the graphical representation of the story map on the PRISM writer interface, the story map creator constructs and keeps track of its data structure. The post processor conducts various tasks: it estimates the minimum and maximum play time, and determines the validity of the story map structure. For instance, the processor determines a story map is invalid if it contains multiple cyclic paths. Finally, the xml encoder converts the story map structure into an xml document—the PRISM container.

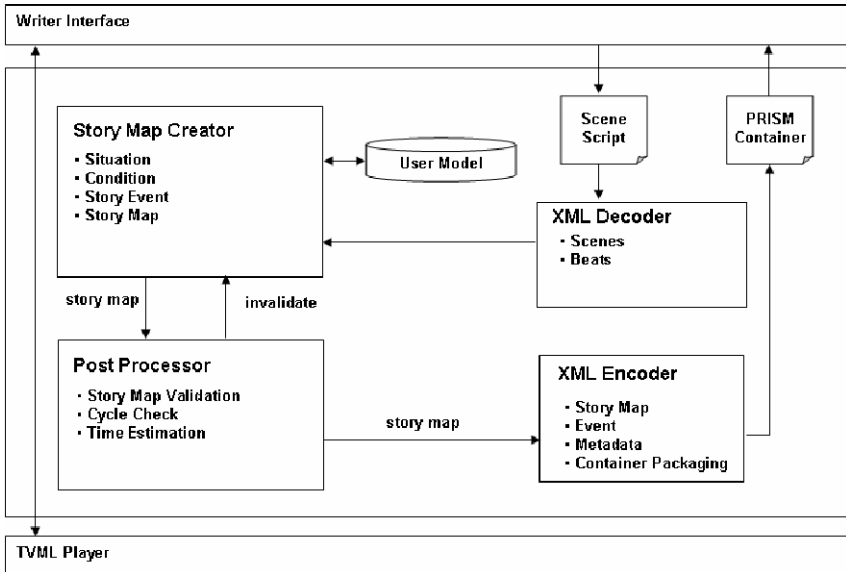


Fig. 4. The PRISM Writer Architecture

4 The PRISM Player

The PRISM Player takes in the PRISM container to realize its TVML scripts on the screen, interacting with the participant. This section discusses the player architecture

and its components, as illustrated in Figure 4. The architecture of the PRISM player focuses on the provision of diverse scenarios to the story participant in real time through her interaction with the story. For this purpose, the PRISM player provides three functionalities: a) decoding a PRISM container, b) managing the interaction between the user and the story, and c) visualizing the story representation.

Upon receiving a PRISM container which is constructed by the PRISM writer, the XML decoder in PRISM parses it to obtain the story information (i.e. story map, story event and metadata) for the story director. The story director then delivers the first situation of the story map to the TVML player, which renders the TVML scripts contained in the situation for the user. The following subsections describe the PRISM player components focusing on the story director and the TVML player, which are two primary modules responsible for interaction management and visual realization respectively.

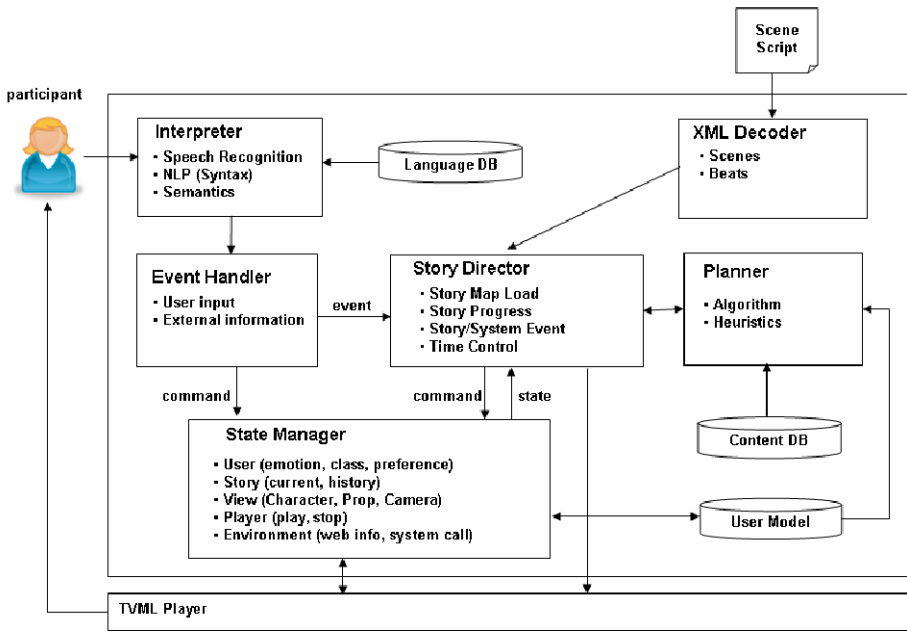


Fig. 5. The PRISM Player Engine Architecture

4.1 The Story Director

The story director takes an interpreted event from the event handler and determines if the event is issued in order to manipulate an item (e.g., characters, props) or to advance the story, where an event can describe the user input, external information, and story world states. For example, the story author makes a character simply wave his hand when the user mouse-clicks on him in a situation, which would not affect the storyline. On the other hand, the user’s decision to take a red pill may lead to another situation, unfolding the story. Events and their treatments are described in the PRISM

container and registered in PRISM during the XML container decoding step so that the story director can refer it to generate reactions to an event.

The user can interact with the system through various input medium such as text, voice, and mouse pointers. To handle the user's natural language interaction, the interpreter processes the following steps: speech recognition, natural language syntactic processing (NLP), and the semantic processing. First, the speech recognition phase converts the user speech input into text. As the second step, our NLP processes the pipeline of tokenization, word tagging, pronoun resolution, named entity tagging, and noun phrase & verb phrase analysis. The final step, the semantic processing, conduct a shallow parsing task that extracts a list of logical sentences from the syntactically analyzed text produced in the second stage, which are then sent to the event handler. When the user input is non-text, the interpreter bypasses the input to the event handler. The event handler encodes the given event as a form of a logical sentence (e.g., *User-TextInput("Help")*, *User-Mouse-Clicked(apple)*) using a simple template matching technique, and then it dispatches the sentence to the story director and to the state manager.

In case where the interpreted event is related to story progression, the story director retrieves the next situation by employing two methods: conditional branching and planning. First, the story director matches the input to the conditions of the current situation. If a matching condition and its branch are found, the story director replaces the current situation with the branch's destination. If no such branches are found, the director delegates its situation retrieving role to the planner. The planner utilizes a planning algorithm—currently a simple forward planning at a depth of one—to find a situation which contains a set of preconditions that are satisfied at the point. This approach has been used as an efficient content selection method in *Façade*, a real-time drama application [8]. However, since content selection based on matching preconditions does not assure the provision of dramatic experiences for the user, the system would require heuristic functions that guide the selection of story elements. For instance, Weyhrauch [19] devises an evaluation function that rates the aesthetic value of the user experiences in an interactive storytelling environment. As such guidance in our system, we plan to devise a generic heuristic function for content selection based on syntactic information of the story graph such as branching factors. On the completion of selecting the next situation, the planner then passes it to the story director.

The state manager maintains various states regarding the virtual story environment and the user. The manager updates the states when an event is input from the event handler or the TVML player. The manager reports a state in response to a query sent from the story director. When events regarding the user model occur, the state manager modifies diverse fields in the user model (e.g. cultural level, educational level, preference, personality, etc.). The user model plays an important role in constructing a personalized story that considers the user's cognitive and emotional states. The user model in the current design employs a planning-based reasoning: a reasoning algorithm, reasoning capacity, knowledge, and preference. As a reasoning algorithm we plan to utilize a partial-order causal planner. To represent the user's knowledge, we will use a set of plan operators which contains preconditions and effects. The user's preference includes heuristic function that controls the content selection processes as described above.

4.2 Visualization

The TVML player [22] visually realizes the story representation into 3D animation. TVML is a text-based language that can describe all the required elements (e.g., stage set, camera, light, characters, prop, voice, sound, etc.) for a 3D animation production. In TVML script, one line of script corresponds to one primitive event. For example, with given a TVML script “character: talk(name = BOB, text = “My name is BOB”),” the TVML Player generates the 3D animation that the character Bob appears on the screen and talks "My name is BOB" with synthetic voice in real time. The TVML script is device independent; thus, it can be thus realized on any platforms (e.g., television sets, mobile devices, personal computers, etc.) where a TVML player is equipped. Our TVML player provides more advanced functionalities than the publicly open source TVML player [22]. The TVML player in our system supplies characters with human-like motions and one-skin model for high-polygon 3D characters. The system also provides a variety of fonts, CG effects such as fade-in and fade-out, and lighting. Furthermore, we plan to provide relative coordination and intelligent interaction among multiple characters and props.

5 Conclusion and Future Work

In authoring interactive narratives, two primary approaches—branching graphs and AI planning—have been significantly used to augment interactivity into conventional linear narratives. Although each approach offers its own possibilities and limitations, few efforts have been made to blend these approaches.

Thus, we have presented a hybrid framework that adapts the conventional branching graphs by integrating AI planning formalism. As a result, the use of branching graphs provides the story author with an intuitive illustration of the story and the story participants with coherent and logical story experience when properly authored. In addition, the use of the planning technique offers the possibility of automated story generation when no predefined branches are available for a given node in the story graph.

Although our initial intention to develop a practical authoring tool for non-experts who need to create interactive stories, it is also expected that this work will be valuable for the interactive storytelling environment systems [7; 15; 20] in which a large volume of high-quality story contents is essential for the provision of varied endings for a given story.

The system is currently under development; we have implemented the content authoring tool and our own TVML player, and designed the xml schema to represent the interactive story structure [10]. We will implement the PRISM player by the end of this year and the PRISM writer by the next year.

Our future work includes the formal evaluation of our system and provision of intelligent assistance for professional knowledge needed work such virtual camera control. Furthermore, we consider employing statistical user models such as Bayesian network. We also plan to supply interface for a variety of user input devices such as the Emotive EPOC™ neuroheadset which detects the human thoughts, expressions, and emotion, in order to maximize the participant’s immersive experiences.

References

1. Carbonaro, M., Cutumisu, M., McNaughton, M., Onuczko, C., Roy, T., Schaeffer, J., Szafron, D., Gillis, S., Kratchmer, S.: Interactive Story Writing in the Classroom: Using Computer Games. In: 2nd International Digital Games Research Conference, Vancouver, Canada, pp. 323–338 (2005)
2. Cavazza, M., Charles, F., Mead, S.: Planning Characters' Behaviour in Interactive Storytelling. *Journal of Visualization and Computer Animation* 13, 121–131 (2002)
3. Donikian, S., Portugal, J.N.: Writing Interactive Fiction Scenarii with DraMachina. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 101–112. Springer, Heidelberg (2004)
4. Iurgel, I.: From Another Point of View: ArtEFact. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 26–35. Springer, Heidelberg (2004)
5. Kelso, M., Weyhrauch, P., Bates, J.: Dramatic Presence. *Journal of Teleoperators and Virtual Environments* 2(1) (2005)
6. Magerko, B.: Story Representation and Interactive Drama. In: The 1st AIIDE, Marina del Rey, CA, USA (2005)
7. Mateas, M., Stern, A.: Façade: An Experiment in Building a Fully-Realized Interactive Drama. In: Game Developer's Conference: Game Design Track, San Jose, CA (2003)
8. Mateas, M., Stern, A.: Structuring Content in the Facade Interactive Drama Architecture. In: The 1st AIIDE, Marina del Rey, CA (2005)
9. Medler, B., Magerko, B.: Scribe: A General Tool for Authoring Interactive Drama. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 139–150. Springer, Heidelberg (2006)
10. Min, W., Shim, E., Kim, Y., Cheong, Y.: Planning-Integrated Story Graph for Interactive Narratives. In: ACM Multimedia 2008 SRMC Workshop, Vancouver, BC (2008)
11. Onuczko, C., Cutumisu, M., Szafron, D., Schaeffer, J., McNaughton, M., Roy, T., Waugh, K., Carbonaro, M., Siegel, J.: A Pattern Catalog For Computer Role Playing Games. In: 5th GameOn North America, Montreal, Canada, pp. 30–38 (2005)
12. Riedl, M., Rowe, P., David, K.E.: Toward Intelligent Support of Authoring Machinima Media Content: Story and Visualization. In: 2nd international conference on Intelligent TEchnologies for interactive entertainment (2008)
13. Riedl, M.O., Stern, A.: Believable Agents and Intelligent Story Adaptation for Interactive Storytelling. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 1–12. Springer, Heidelberg (2006)
14. Riedl, M., Young, R.M.: From Linear Story Generation to Branching Story Graphs. *The IEEE Journal of Computer Graphics and Applications*, 23–31 (2006)
15. Riedl, M.O., Stern, A.: Believable Agents and Intelligent Scenario Direction for Social and Cultural Leadership Training. In: Behavior Representation in Modeling and Simulation Conference, Baltimore, Maryland (2006)
16. Spierling, U., Weiß, S.A., Müller, W.: Towards Accessible Authoring Tools for Interactive Storytelling. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 169–180. Springer, Heidelberg (2006)
17. Swartout, W., Hill, R., Gratch, J., Johnson, W.L., Kyriakakis, C., LaBore, C., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Morie, J., Rickel, J., Thiébaux, M., Tuch, L., Whitney, R., Douglas, J.: Toward the Holodeck: Integrating Graphics, Sound, Character and Story. In: 5th International Conference on Autonomous Agents, Montreal, Quebec, Canada, pp. 409–416 (2001)

18. Tully, G., Turner, S.: Integrated Decision Points for Interactive Movies. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) TIDSE 2004. LNCS, vol. 3105, pp. 61–67. Springer, Heidelberg (2004)
19. Weyhrauch, P.: Guiding Interactive Ficiton. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA (1997)
20. Young, R.M., Riedl, M.O., Branly, M., Jhala, A., Martin, R.J., Saretto, C.J.: An Architecture for Integrating Plan-Based Behavior Generation with Interactive Game Environments. *Journal of Game Development* 1(1), 51–70 (2004)
21. INSCAPE Storytelling, <http://www.inscapers.com/>
22. TVML, <http://www.nhk.or.jp/str1/tvml/english/player2/index.html>

Tales for the Many: Process and Authorial Control in Multi-player Role-Playing Games

Anders Tychsen

Center for Computer Games Research, IT University of Copenhagen,
Rued Langgaards Vej 7, 2300 Copenhagen S, Denmark
anty@itu.dk

Abstract. Multi-player Role-Playing Games form rare examples of interactive, collaborative storytelling systems, especially within the tabletop format. However, despite recent advances in the theory and general knowledge of the field, models of the operation of the storytelling process in multi-player non-digital role-playing games remains relatively unspecific. In this paper, a critical synthesis of the available theory is combined with observations from a series of experiments, in a detailed model for the gaming process.

Keywords: Role-Playing Game, collaborative storytelling, game master, game process.

1 Introduction: Non-digital RPGs and Storytelling

With the increased focus in the games industry on reaching casual and current non-gaming segments of the market, and the interest in interactive storytelling [12,13,16,17] and agent-based systems [20], the ability to provide personalized experiences to users is becoming increasingly important. However, technological and resource limitations mean that e.g. game designers are forced to create the stories in virtual worlds with technologically and financially imposed limits on player freedom [2], for example by relying on pre-authored plots [7]. Table-top Role-Playing Games (RPGs) (also called Pen-and-Paper RPGs) are therefore subjects of increasing interest in academia and industry as a source of inspiration in the development of interactive storytelling systems, for use in computer games, education and interactive entertainment [12,13,20]. While table-top RPGs are expressed through a different medium than computer-based RPGs, the game forms share several features, notably the underlying rules systems and mechanics, thematic game worlds, incentive systems and often a lack of clearly defined victory conditions [10]. Even though table-top theory has taken great steps forward in recent years, there has yet to be successful attempts at modeling these games in a manner that is conducive to the creation of a digital variant with the same in-built flexibility. This has several reasons, notably that much of the theory e.g. [3,5,8,11,14,15,18] has not been created with this purpose in mind but to advance the design and gaming experience of RPGs [24], or to investigate RPG culture [6]; secondly that the game process itself is reliant on human interaction, which spans a variety of responses and cognitive processes that are challenging to

model e.g. [12,13,17]. A line of thinking has emerged in the past few years that may eventually lead to models of the RPG gaming process which can be utilized in the construction of digital systems [1,4]. Kim [10], inspired by discussions on online discussion forums, developed a model of the RPG gaming process that outlined communication channels between the game participants and integrated the concept of a shared game space where the players communicated the actions of their characters within the shared, imagined, fictional world. Henry [8] advanced this idea by modeling the basic information flow of RPGs, with the game components and participants forming a network of data sources and entities. Mäkelä et al. [13] further addressed this line of thinking by deconstructing RPGs into a series of processes. Tychsen [21] noted that RPGs could be modeled as information systems, combining the previous ideas into a coherent framework.

This paper presents a critical synthesis of the existing models of the process of non-digital RPGs, and theory pertinent to the actual collaborative storytelling process in multi-player tabletop RPGs; and presents a model of the RPG process from an information systems viewpoint. The model is based on existing knowledge and in part on observations during a three-year project on storytelling and player interactions in multi-player RPGs across non-digital and digital platforms. During the project, more than 150 hours of multi-player RPG sessions were run and the data sources include audiovisual recording, observation, screen capture (of digital RPGs), surveys and interviews [23]. Focus will be on non-digital RPGs - the way stories are formed between the game participants, the distribution of authorial control and the operations of the Game Master (GM), a participant who traditionally is responsible for directing the gaming experience and provide feedback to the players as to the effect of the actions of their fictional characters in the context of the game world. The GM is therefore an analogy of an automated interactive storytelling system, although the interaction between the players is also important to the gaming experience [1,4].

2 Building Premise

When a group of persons engages in a role playing game, there is a range of conditions that must be met before playing can begin. These conditions describe the knowledge that the players must have about the game in order to play it, what expectations the players have from the game, how they are going to be playing it, genre, characters, distribution of authorial control, style of play etc. The establishment of these conditions by the players and the agreement or framework for the game thereby established is termed the game **premise**. If the players do not agree to a premise, their different assumptions, expectations and goals will conflict during play. Premise can be broken down into components, all of which are renegotiable during play:

1) Fictional game world: When embarking on a RPG, the players have to agree upon a specific fictional game world that forms the setting, or stage, for the game. The game world can be specified to great detail, including such features as a history of the game world, geological/biological information on its nature and appearance, any civilizations and so forth. The game world needs to be understood to avoid confusion.

2) Fictional contract: Fictional worlds are governed by rules, which determine how the fictional world functions, e.g. the damage of specific weapons, whether the laws of physics are observed, and so forth. Some of these rules impact the players directly and is usually quantitative (the damage of weapons), some indirectly (gravity means you need an airplane or magic to fly). Because the fictional world has to follow certain rules, this means that the fictional world conversely can affect the game. This leads to an important concept, the fictional contract: *The agreement between the players as to how the world fiction operates and what rules govern it*. The fictional contract ensures that even though every player may visualize the fictional world setting differently, they agree as to how their characters interact with the fiction, and the basic principles of the fictional world. The fictional contract is integral in facilitating the understanding of the fictional world itself and how it operates. Without it, rules and setting are not communicated successfully to the players. Rules of communication, authorial control etc. between the game participants and similar mechanics that occur outside the fictional world does not fall under the fictional contract, however, it does detail how the *characters* can interact.

3) Player characters: Each player controls usually one, but potentially more (e.g. in the RPG *Ars Magica*) fictional characters, which are generally the protagonists of the game story. In relation to the establishing the premise, these characters must be defined in terms of their stats, abilities, personalities etc. to the extent desired by the players and/or required by the game rules.

4) Visuals: While not a strict requirement of the premise, illustrations, music, sound tapestries, and similar game elements can be utilized to enhance the atmosphere and visualize the fictional world or similarly. In a digital RPG, color is often provided via an introduction sequence, which also serves to communicate and reinforce the fictional contract (e.g. *World of Warcraft*, *Neverwinter Nights*, *Baldur's Gate*). The general visual element of digital RPGs also serve in this regard and potentially has an advantage over non-digital RPGs in ensuring all participants have a common understanding of the game world, which can be a problem in RPGs [20].

5) System: Because multi-player RPGs are based on a shared understanding of events occurring in an imagined, fictional game world, the method for which these events, the game world and the associated character actions are described and agreed upon is vital to the game process. Edwards [5] and Kim [11] referred to the shared understanding of the fictional game world and the events taking place therein as the “shared imagined space”. Importantly, while the fictional contract governs the functionality of the game world and provides the framework and game rules for character actions, it is somewhat malleable and subordinate to the principles governing the actual game play, or the game system. The system of a RPG is the means by which the participants distribute the authorial control, or credibility, to make statements about what takes place within the shared imagined reality of the game world (the shared imagined space) [24].

Traditionally, players have the credibility to define the actions of their characters, and what their characters say. The GM (who can have various titles), has the credibility to decide on the success or failure of actions, possibly guided by the game rules, and the reaction of the game world and the inhabitants of the world, to the

actions of the characters – generally guided by the need for game story, entertainment value etc. Authorial control is important to the development of the collaborative story, as it determines how the participants can affect the game world, the other characters, and thereby the game story. Importantly, system is not the same as game rules. The game rules may define that a sword does more damage than a knife; however, it is the system that defines who has the credibility to use the knife in the first place. In essence, rules have authority to give credibility to statements; however, rules do not have authorial control in themselves [25]. The players need to decide whether a rule applies, and how to interpret it, before the authority of the rule is valid. The use of dice and other resolution mechanics are like rules a means of supporting the game system, providing authority. E.g. a player can roll a dice to determine whether a given character action was successful according to the rules. If this succeeds the statement of the player that the action was taken has substantial credibility. Whether a GM or the other players can ignore the rules is also part of the game system.

6) Social contract: While the rules and the system govern the interactions between the players and the characters, the characters and/or the fictional world, and objects/entities within the fictional world, they do not govern interactions between the players. The social contract is the umbrella term utilized for the rules of social conduct during game play, the ambitions and desires of the players with the game playing activity, and what style of play that to be carried out. The agreement to the above conditions bears resemblance to the concept of *lusory attitude* [19], which is defined as a certain state of mind that is required of players for them to enter a game. The formation of the game premise can be an extended process, especially if a group of strangers gather around a game table to play. While running the game experiments mentioned in the introduction [23] it was observed that the premise-building phase could take an hour or more to accomplish in these cases. In the groups observed where the players had prior joint gaming experience, this process was reduced in temporal length. Following an introduction to the rules system used, the game characters and the fictional world, these groups quickly began playing.

3 The Gaming Process

The gaming process of a multi-player RPG can be viewed from a variety of angles, and this has caused substantial problems in defining a coherent model for these games. There is no one correct way of modeling these games, and different theorists have applied different methods in order to evaluate specific features. The approach that is adopted here, which aims at incorporating existing theory with observations from the empirical work described above, is to separate what could be termed the “**processing**” component of RPGs from the “**storytelling**” components. The processing component covers the basic framework of these games, modeled and described as an **information system** [26]. The way the participants handle authorial control, conflict resolution and other elements of the collaborative story can vary, however, importantly they will always utilize the same communication channels, and will have the same means of processing, storing and retrieving information.

3.1 RPGs as Information Systems

Pen-and-Paper RPGs are games where the state of the game world changes due to the actions of the fictional characters within the game world, as directed by the players. A cyclic behaviour in the system originates because there exists a communication between the real world of the players, and the fictional game world of the characters. The decisions of the players lead to their characters taking specific actions. Via the GM, the game world (or objects/entities therein) provides a reaction, which is fed back to the players, who subsequently processes the feedback, before making a decision about what to do next (individually or as a group) (Figure 1). With each cycle (or each step of the cycle), the game is brought to a new state. This top-down view of RPGs is applicable to all forms of these games and encompasses the interactive nature of the game playing activity itself. This behavior can be further detailed by considering RPGs as information systems. In this context, an information system is a collection of people, processes and technologies (in this case books, dice, rules, possibly computers, artwork etc.) to support the information needs of the participants and provide the framework for the gaming activity. In this viewpoint, the game elements and the gaming process itself is described in terms of distinct, dynamic and static data sources, entities that actively interact via specific communication channels and which support the processing of information, with the purpose of generation the information to support the game and the game story.

Different entities have different information needs, and different rights of data access and ways of affecting information flow and -processing. In the context of RPGs, these rights depend on the distribution of authorial control. For example, a GM might have different information requirements than the players, e.g. in relation to knowing the motivations of NPCs, and will also have the rights to access additional information and change it in real-time. The overall action-reaction system of RPGs represents a simplistic version of the actual gaming process that takes place in these games. While the generally cyclic nature of progress from game state to game state appears to be valid, each cycle can be subdivided into numerous component processes, depending on the level of detail required in an analysis.

During the observation of the RPG experiments mentioned above, it was commonly observed that the participants were able to follow and/or participate in multiple conversations at the same time, jumping back and forth in the game story, in the chronology of the game world, and even following different events in different geographical locations of the imagined, fictional world. It was not unusual to observe individual players conversing with two or three other players during tense sections of the game sessions, producing a sound landscape that is incomprehensible to any observers, while being deeply immersive to the participating players.

In order to be able to model this level of detail, the feedback cycle of RPGs need to be broken down into systems of sub-processes (Figure 2). A sub-process can occur at any of the stages in the regular feedback cycle (decision – action – reaction – processing). Importantly, players of RPGs have the opportunity of processing information (decide on their response to the game world state update), before their characters take action. This processing stage can range from the very simple to the complex. Players do not react uniformly to a given game world state change, and can even react at different times to the same input from the game. During the processing

stage, more than one sub-process can operate at the same time. For example, in a multi-player RPG, a game world state update from the GM can result in the players internally discussing what their next action should be and queries to the GM for detailed information about a specific section, object or entity of the game world. Furthermore, players need not make their characters react at the same time in the same way: One player can engage opponents while another talks to a merchant NPC. In this case, the cyclic process is split into several sub-systems, up to one for each individual player. These sub-processes keep operating until a point in the playing of the game where the player characters again move or act more or less coherently, e.g. by the GM aligning all the characters at the same geographical, chronological and contextual point within the game world and game story.

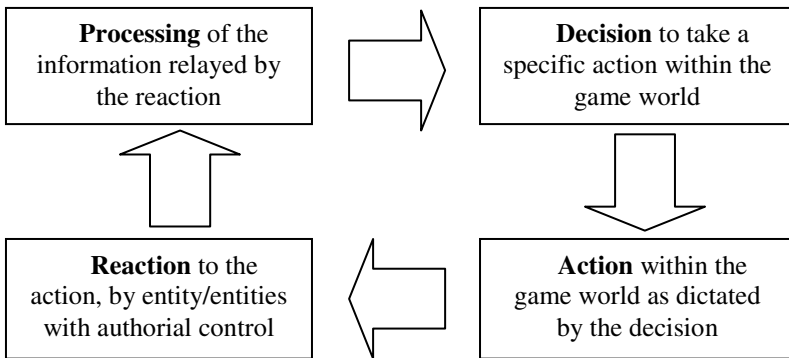


Fig. 1. A simplified view of the information feedback cycle of RPG gameplay

What is obvious from the above is that communication and actions in RPGs can be viewed at different scales of resolution. Importantly, game progress equals a change in the game world state [9], and how this is defined depends on the scale of resolution. The cyclic model can be described at various levels of resolution or with the focus aimed at different forms of game state progress (narrative progress, advancement in the tasks the game sets for the players, advancement in game mechanics, e.g. character level). In a digital RPG, moving an avatar from point A to B causes a game state change. However, such an action will not necessarily cause communication, result in narrative progression or similar. Similar to non-digital RPGs, actions, player communication etc. can be viewed at different scales of resolution. Likewise, sub-processes can be of different scale, depending on the level of resolution of the main process model (Figure 2). If a sub-process is identified as causing substantial game progress, the feedback loop in question is a double loop, and needs to be separated into two serially connected cycles. Defining game progress is here purposefully left vaguely defined as the variable resolution value of the interaction model presented here means that game progress will need to be specified for each specific application of the model.

An important feature of the RPG information systems is that the entities have varying degrees of data access to a series of dynamic and static repositories [21]. In a

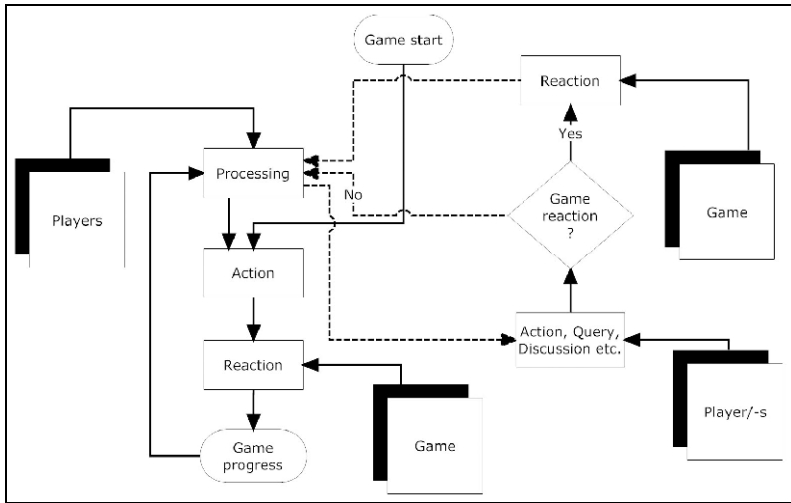


Fig. 2. A model of the game process of RPGs (left), including a sub-processing loop (right). From a starting condition, players act upon the game world via their characters, prompting a response and update of the game world state, as decided upon by the GM. This leads to progress, and the process repeats. Symbols as for standard data flow diagrams (DFDs).

RPG, the players typically have access to their character, rules, and some knowledge of the fictional world, e.g. through books provided by the game publisher. The character can be viewed as a dynamic repository, as it will change status throughout game play, e.g. it may become wounded or gain new abilities. Rules are in comparison static; however, as noted above the participants have the option of changing the rules of a RPG. The GM, normally has access to additional sources, notably the game module (the manual for the story), as well as additional game world information. In a digital RPG (computer-RPG or Massively Multiplayer Online RPG, MMORPG) the game code/engine would feature as an additional set of information that can be accessed. Related to the subject of data access is content creation rights and regulations; however, this is a subject that is out of scope of the current study.

3.2 Creating the Game Story

Modeling RPGs in their capacity as information systems provides knowledge about how the game process operates, however not how the authorial control operates in shaping the game story. In order to analyze this process, the basic concept of authorial control has to be expanded. The game story of a RPG can be defined as the sum of everything that takes place within the fictional world (this definition purposefully ignores the debate about games, stories and interactive narratives). This does not include the communication between the players that does not relate directly to the actions of their characters – for example, questions about the game rules or requests for information [22]. When observed from the outside, the story of a RPG appears to run progressively from the opening scene to the final scene, and it is not immediately obvious that the players and GM are continually selecting between different paths as

they go along. [22] divided the game story of a RPG into **perceived** and **conceived** parts, with the former being the part of the game story that has been communicated through the shared imagined space (the shared play space).

The conceived story rests within the pre-planned materials and notes of the GM, as well as the mind of the GM and the players, i.e. their intentions for the game story and the un-mapped actions of the characters. As a result of the cyclical game process, the conceived story is gradually transformed to a specific perceived story. The process through which this transformation takes place is at least partly cognitive, and has yet to be analyzed in detail; however, it can to some degree be described by analyzing the distribution of authorial control and the way that the GM manages the game story.

Variations in authorial control, different means of conflict resolution and styles of GM story management in conjunction provides a very wide solution space that RPG groups can adopt in order to facilitate collaborative storytelling in a manner that suits their specific ambitions: Storytelling in RPGs can vary from the GM and the players establishing an initial situation in the game world, and improvising from there, making up the story and its characters as they go along; to a strictly linear experience where the players merely fill in the blank space in a rigid story structure.

Compared to digital RPGs (single-, multi- and massively multi-player), such as *Neverwinter Nights*, *Temple of Elemental Evil*, *World of Warcraft*, *Lineage*, the variety of storytelling modes in table-top RPGs is greater, because digital games are constrained technically and financially. Technically, a human GM has the ability to on-the-spot create new environments and situations, while a digital game requires pre-programmed content, although it should be noted that techniques for procedurally generated and emergent content in digital games holds the potential to expand the flexibility of digital RPGs. Financially, while a GM can invent content for free, developing content in digital games costs resources, and any game production is limited in this regard. From the perspective of the player, table-top RPGs provide the potential to direct player characters to perform any plausible action within the imagined fictional environment, unlike in digital RPGs where the numbers of choices are strictly limited by programming. Adding multiple players in the digital RPG context somewhat remedies the situation by adding unpredictability; however, the basic technological framework of these games means that there are limits to the flexibility possible in terms of storytelling. These limitations have not prevented developers from creating immersive and engaging, multi-linear and/or branching storylines and characters in digital games, e.g. *Knights of the Old Republic*, *System Shock II*, *Bioshock*, where the game engine itself takes on a role akin to that of the GM. However, the operational space is smaller compared to the shared imagined space of tabletop RPGs. This is also the case in MMORPGs, where role-playing oriented guilds in e.g. *World of Warcraft* have managed to transfer some of the storytelling principles of table-top RPGs and live action RPGs, and adapt them to the unique context of MMORPGs (i.e. chat-based interaction, use of props and costumes), but which are similarly limited in the flexibility of what they can show and do within the confines of the virtual environment.

Personal vs. shared story: Before considering the story generation process in detail, it is important to realize that there is no single game story arising from a RPG game session. Rather, a series of stories that vary depending on the eye of the beholder: First of all, there is the game story that can be recorded and observed in a RPG

gaming session. Secondly, each participant will have formed their own version of the game story. This second viewpoint is related to a key feature of RPGs, the concept of a personal imagined space, which supplements the shared imagined space. The personal imagined space operates within the mind of the individual gaming participants, and it is here that the personal perception of the game story is being formed, which can vary to a greater or lesser degree from the story expressed in the shared imagined space and the personal imagined spaces of the other game participants [11]. While this has not been examined empirically, anecdotal evidence indicates that the variations in the personal imagined game stories are largest at the fine level of story detail, e.g. most players will imagine the characters in a room if the GM informs them they are in one, even though the image of the room conjured will likely vary from player to player. Where the story expressed in the shared imagined space can be recorded directly, the personal story can only be accessed indirectly, i.e. via the player recounting their perception of the game. For simplicity, when discussing game story in the following, the shared and personal stories of the game participants are evaluated as a coherent construct.

Distribution of authorial control: As an entity in a game system, the traditional role of the GM is to provide information about the fictional game world and the development of the storyline (to the best of his ability, considering that the players are, at least to some degree, in control of the story protagonists), refine it through player queries, and arbitrate player interaction with the rules structure, while following the conditions of the premise and ensuring that the various ambitions of the players (and the GM) with regards to the gaming activity, are met. The players in turn utilize the input and description from the GM to formulate a response and new actions, and to construct individual mental model of what events are taking place. In other words, the GM is in charge of facilitating game flow and –story, provide the environmental content of the fictional reality, and possibly arbitrate conflicts. The GM is therefore a participant but rarely a player. The degree of authorial control the GM has depends on the level of credibility of the players (a RPG could even operate without a GM, by distributing authorial control completely among the players): The more authorial control the players have, the more adaptable and flexible the GM has to be in the management of the game story. An extreme case is represented by most digital RPGs, where the players have a very limited authorial control beyond opening boxes, attacking MOBs and carrying out pre-planned conversations with NPCs. The distribution of authorial control among the players is related to the concept of **stance** [5]. Stance refers to the relationship between the character, the player and the shared imagined space and determines: 1) The degree to which the player can direct the character according to his/her own motivations, or has to follow the motivations of the character; 2) The degree to which the player can affect the fictional game world in a manner that is outside the abilities of the character (Table 1), The traditional GM can be viewed as an extreme case of the Director stance, which is unbound by a specific character. While authorial control can be described and defined for any given RPG, the distribution of authorial control in a game session varies during game play. Players can change stances with regards to their character, the GM can vary the amount of control imposed on the game world, etc., and the game play can therefore be compared to a continuing negotiation process where different participants discuss, debate and propose statements about events occurring in a fictional game world.

Table 1. Definition of four types of stance (source: Modified from Young [25])

Stance	Definition
Pawn	The character is a token utilized by the player to act within the game world. The character performs only actions that the player wants it to do. It does not matter if the actions of the character make sense within the world fiction. The player is limited to interact with the game world via the character.
Actor	The player directs the actions of the character according to what the player believes the character would do in the given circumstances. Authorial control of the player is limited by the character as well as the thoughts of the character and the abilities it has to affect the game world.
Author	The player can utilize his/her own knowledge and motivations in deciding on character actions, and thereby apply direct control on the game story. However, character actions need to be justified in terms of the character, e.g. by creating a reason for why the character acted in a given way.
Director	The player has the authorial control to directly affect the shared imagined space (the game world), which is outside of the control of the character. The degree to which this control can be applied can vary substantially.

The illusion of authorial control: The division of and fluctuations in authorial control in a RPG situation generally determines how the collaborative story is formed at the point between the conceived and the perceived story. The story management of the GM is in turn directly dependent on the way authorial control is distributed. Different distribution models lead to specific ways of managing the game story. Importantly, because authorial control can fluctuate throughout a game session, the style story management can also vary. In the above, it has been assumed that authorial control was real. However, the authorial control that a player has to affect the game world is not necessarily directly proportional with the amount of control the player has over the game story. E.g., a player may direct a character to perform an action that eliminates an NPC, and have the credibility to do so. However, the GM may subsequently decide to resurrect the opponent NPC, or introduce a different NPC with similar importance to the game story. The player may not ever realize that his/her action did not result in a substantial alteration of the game story, merely a change in the delivery of it. In this situation, the player is subjected to an illusion of authorial control, a perceived freedom that does not in fact exist. A simple way of illustration how illusionary authorial control can be used by GMs (and game designers) is to consider the situation where a player character is positioned in a room with two doors. Irrespective which door the character chooses, the GM decides that it will lead on into the same room. Should the player check first one door, then another, the GM can in real-time improvise a means of preventing this behavior, e.g. by having a noise occurring from the door the character checks first, prompting the player controlling the character to turn back to the first door. The player may believe that the character had freedom to choose both doors, without this being the case.

The relationship between perceived and actual authorial control is important to GM management styles: In one extreme, the actions of the player characters are devoid of impact to any significant degree. The GM counters and adapts the game story to any action the characters may take, and the players are not aware of this. The GM is at all times in control of how the story will develop and the level of variance permitted to

the players. It may also be the case that the players are aware that the GM is in effect telling them a story, and that they can only affect the game story in a very limited capacity. The GM is in this case still utilizing the illusion of authorial control; however the players accept it. This form of story management is comparable to digital RPGs where technical and financial limitations impose limits on flexibility [2,7].

With increasing player freedom, decisions will impact on the game story beyond the insignificant. The story is adjusted throughout play to accommodate the actions of the players, however, the players implicitly or explicitly follow the clues, prompts and hints of the GM. The challenge in this form of RPG is the completion itself; however, the path from beginning to end can vary. The game module utilized by GMs can vary from defining a sequence of scenes/events, a network of events or a list of conditions and circumstances. The smaller the degree of pre-definition of the game story, the more challenging it is for the GM to keep the game story under his/her control, and the more the players will be able to impact the story. Further increasing player freedom reduces the GM to establishing the fictional game world and possibly some key situations or NPCs, but otherwise falls back to a role where he/she reacts to the actions of the players and attempts to improvise a story. This type of RPG play requires high engagement on behalf of the players, as well as GM management. The players enjoy a high degree of authorial control in the actions of their characters [25].

4 Conclusions

In this paper a critical synthesis of current role playing game theory combined with observations from a series of empirical experiments with table-top RPGs [22,23]. The preparational steps and mechanics of the gaming process have been modeled, and the division of authorial control utilized by game participants, including the GM, to create collaborative storytelling environments has been examined. This study and those referenced herein represent however initial steps in the way to fully analyze the RPG situation and ideally create heuristics for digital storytelling systems [12,13,17,20]. Because RPGs are games that operate partly within the mind of the participants, it remains uncertain if the principles of GM behavior can be transferred directly to digital systems, or whether only partial models can be integrated. Future studies in this area will therefore need to include analysis of the cognitive process of story management that takes within the mind of GMs – as well as players. When interviewing the GMs involved in the empirical multi-player RPG experiments, as to how they manage RPG games, a wide variety of answers were provided, and most agreed that the process was partly improvised: Because RPGs are based on conversations in real-time, the GM will at some point have to improvise or role-play on-the-spot, without the ability to consult NPC charts, scene overviews and similar. There was further agreement that the GM during a game session continually revises the ideas or plans for the game story, at several different levels of story detail and planning, based on the interactions with the players.

Acknowledgements. Numerous discussions with colleagues have been seminal in forming the ideas presented: Dr. Janet Murray, Dr. Michael Hitchens, Dr. Ruth Aylett, Dr. Sandy Louchart, the DiGRA RPG SIG, the *Forge* and *RPG.net* communities.

References

1. Aylett, R., Louchart, S., Tychsen, A., Hitchens, M., Figueiredo, R.: Managing emergent character-based narrative. In: Proceedings of INTETAIN 2008, Cancun (2008)
2. Bartle, R.: *Designing Virtual Worlds*. New Riders Games, Berkeley (2003)
3. Bøckman, P., Hutchison, R. (eds.): *Dissecting larp*. Grimshei trykkeri AS, Tampere (2005)
4. Combs, N.: The Intelligence in the MMOG: From scripts to stories to directorial AI. In: Proceedings of the OP Conference, Copenhagen, Denmark (2004)
5. Edwards, R.: *GNS and Other Matters of Role-playing Theory*. Adept Press (2001)
6. Fine, G.A.: *Shared Fantasy: Role Playing Games as Social Worlds*. University of Chicago Press, Chicago (2002)
7. Hallford, N., Hallford, J.: *Swords & Circuitry*. Prima Tech, Roseville (2001)
8. Henry, L.: *Group Narration: Power, Information, and Play in Role Playing Games* (2003), <http://www.darkshire.net/~jHKim/rpg/theory/liz-paper-2003/>
9. Juul, J.: *Half-Real: Video Games between Real Rules and Fictional Worlds*. MIT Press, Cambridge (2005)
10. Kim, J.H.: *Story and Narrative Paradigms in Role-Playing Games* (2003), <http://www.darkshire.net/jHKim/rpg/theory/narrative/paradigms.html>
11. Louchart, S., Aylett, R.: Solving the narrative paradox in VEs – lessons from RPGs. In: Rist, T., Aylett, R.S., Ballin, D., Rickel, J. (eds.) *IVA 2003. LNCS (LNAI)*, vol. 2792, pp. 244–248. Springer, Heidelberg (2003)
12. Louchart, S., Aylett, R.: Emergent Narrative, requirements and high-level architecture. In: Proceedings of SETN 2004, Pythagorion, Samos (2004)
13. Mäkelä, E., Koistinen, S.: The Process Model of Role Playing. In: Bøckman, P., Hutchison, R. (eds.) *Dissecting larp*, Knutepunkt, Tampere (2005)
14. Mackay, D.: *The Fantasy Role-Playing Game: A New Performing Art*. McFarland and Company (2001)
15. Murray, J.: *Hamlet on the Holodeck*. Free Press, New York (1997)
16. Peinado, F., Gervás, P.: Transferring Game Mastering Laws to Interactive Digital Storytelling. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) *TIDSE 2004. Incs*, vol. 3105, pp. 48–54. Springer, Heidelberg (2004)
17. Salen, K., Zimmerman, E.: *Rules of play*. MIT Press, Cambridge (2003)
18. Suits, B.H.: *Grasshopper: Games, Life, and Utopia*. University of Toronto Press (1978)
19. Trappl, R., Petta, P. (eds.): *Creating Personalities for Synthetic Actors: Towards Autonomous Personality Agents*. Springer, Berlin (1997)
20. Tychsen, A.: Role Playing Games – Analysis across Two Media Platforms. In: Proceedings of IE 2006, Sydney, Australia, pp. 75–82 (2006)
21. Tychsen, A., Heide-Smith, J., Hitchens, M., Tosca, S.: Communication in Multi-player Role Playing Games - The Effect of Medium. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) *TIDSE 2006. LNCS*, vol. 4326, pp. 277–288. Springer, Heidelberg (2006)
22. Tychsen, A., Newman, K., Brolund, T., Hitchens, M.: Cross-format analysis of the gaming experience in multi-player role playing games. In: *DIGRA 2007 Proceedings*, Tokyo, Japan, pp. 49–58 (2007)
23. Young, J.M.: Theory 101: System and the Shared Imagined Space. *Places to GoPeople to Be*, 26 (2005a), <http://ptgptb.org/0026/theory101-01.html>
24. Young, M.J.: Theory 101: The Impossible Thing Before Breakfast. *Places to Go, People to Be*, 27 (2005b), <http://ptgptb.org/0027/theory101-02.html>
25. Whitten, J.L., Bentley, L.D., Dittman, K.C.: *Systems Analysis and Design Methods*. McGraw-Hill, New York (2003)

An Intelligent Plot-Centric Interface for Mastering Computer Role-Playing Games

Carlos León, Federico Peinado, Álvaro Navarro, and Héctor Cortiguera

Departamento de Ingeniería del Software e Inteligencia Artificial
Universidad Complutense de Madrid, Spain
{cleon,fpeinado}@fdi.ucm.es,
{alvaro.niborra,h.CORTIGUERA}@gmail.com

Abstract. Role-Playing Game Mastering is one of the most cited examples of “ordinary life Interactive Storytelling” in scientific literature. There are some computer interfaces for game mastering, but most of them are designed as part of the toolkits that are used to control low-level events of specific videogame engines. In this paper a more abstract application has been developed, an intelligent graphic interface that allows human Game Masters to direct stories in virtual environments from a more comfortable and narration-oriented point of view.

1 Introduction

The Game Master (GM) of Role-Playing Games (RPGs) is the most intuitive model of Interactive Storytelling found in ordinary life [1,2,3,4]. Despite this fact and the need of high-level control of game engines [5], there is little reported work on implementing this approach to direction of narrative virtual environments.

Due to the complexity and difficulty of adapting game contents to the unpredicted actions of the players, current approaches to mastering Computer Role-Playing Games (CRPGs) usually depends on particular game systems. This forces GMs to learn how each particular system works internally and how it should be controlled using a specific Graphic User Interface (GUI).

This paper presents a working prototype of an engine-independent plot-centric GUI for game mastering. Its goal is to improve the toolkits available to GMs for directing interactive storytelling in a virtual environment.

2 A Narrative Approach for Role Playing Mastering

In order to develop our proposed GUI for game mastering, a layered architecture is used, which hides details in high-level *layers* allowing a more narrative-like managing of the gameplay. The architecture, described in Figure 1 is the ground on top of which we create the real system, called IMP.

In the main window of our IMP everything can be modified by using the mouse, dragging boxes and arrows. It is possible to create whole stories from scratch by

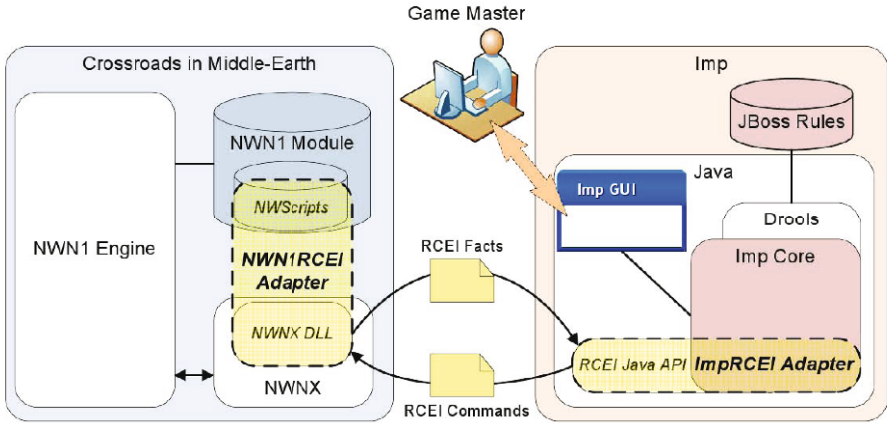


Fig. 1. The layered architecture of IMP

adding characters, setting their characteristics, creating events, etc. The main window has three areas: The Timeline, the Objects Tree and the Characteristics Editor.

The Timeline contains a schematic, time-ordered representation of the game story. The GM can easily edit its events: Moving them by *dragging and dropping* with the mouse, creating new ones or deleting them. IMP displays a layout where characters' timelines are presented and, for each character, its sequence of events are laid out from left to right.

Interactivity forces the application to consider some time-specific notation in the interface. The GUI for this system is provided with a *time-bar*, which moves from left to right in the Timeline area showing the current time in the gameplay. With the time-bar, GMs have information about the moments their events or relations are going to be executed in the system.

The Objects Tree shows a hierarchical layout of the narrative objects inside the game. These objects can be *characters*, *places* or *items*. Relations between these narrative objects are also present as elements in this tree. Using the Objects Tree, the user can easily access to any object of the game. In order to configure characters, locations or events, the Characteristics Editor allows the GM to set the details for each object.

Monitoring the game state is performed without user intervention. IMP receives commands from the game and updates the main window information that layouts automatically. New events arise in the Timeline area, and character's properties change as the game continues. Then, the GM can keep on modifying it or just watching the story as it is created by the players' actions. Stories can be saved in a human readable XML format and then, be reloaded in IMP for a later study of the game, or to replay a version of the story.

The system operates according to the next sequence: after connecting IMP with the game engine, it creates messages corresponding to the current events being played (those which are under the time-bar), and sends them to the game.

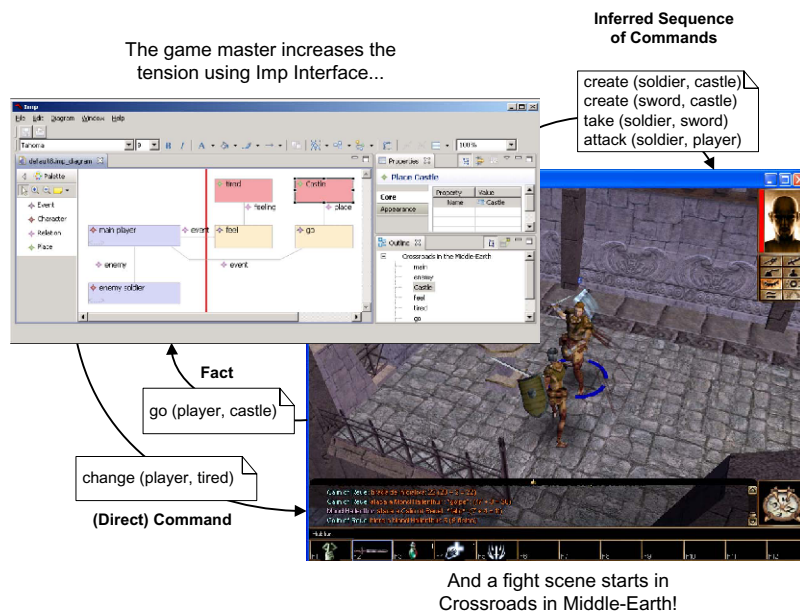


Fig. 2. A Game Master using IMP for directing a RPG

Once the game engine run those messages, new information is created inside the game and sent it back. IMP, then, gets that information, inferring narrative facts that are not explicitly represented inside the game and updating what is shown to the GM.

An example of how IMP is used for directing a RPG is shown in Figure 2.

The engine of *Neverwinter NightsTM* (NWN) is the technology we have chosen for implementing our test game: “Crossroads in Middle-Earth”. NWN is a CRPG with that present opportunities for high-level creation of narrative scenarios [6]. IMP is connected to the NWN engine via RCEI (Remote-Controlled Environments Interface [7]) and *Neverwinter Nights Extender* (NWNX), which offers a basic but powerful communication pipe between our application and the game, getting results without too much development effort. Finally, *JBoss Drools* [8] has been chosen as the rule-based system behind our application, which is fully integrated with Java, the platform in which IMP is implemented.

3 Conclusions

In this paper a new plot-centric GUI for game mastering called IMP is presented. This application is designed to improve GM performance in the direction of CRPGs. Firstly, GMs will not need to learn the secrets of many particular game engines, focusing solely on gameplay after learning how to use a simpler GUI. Secondly, they will benefit from the efficiency of high-level direction instead of dealing with low-level details of the virtual environment. Thirdly, rules for game

management are defined in a editable file so the GM can customize the behaviour of the system. Nevertheless it is true that there is a loss of control granularity due to the abstraction effort required for communicating two different systems in an independent way.

The next step is to evaluate IMP in game situations in order to probe that the game experience has been significantly improved.

Acknowledgements

This research is funded by Ministerio de Educación y Ciencia (TIN2006-14433-C02-01), Universidad Complutense de Madrid and Dirección General de Universidades e Investigación de la Comunidad de Madrid (CCG07-UCM/TIC 2803).

References

1. Louchart, S., Aylett, R.: Solving the Narrative Paradox in VEs: Lessons from RPGs (2003)
2. Peinado, F., Gervás, P.: Automatic Direction of Interactive Storytelling: Formalizing the Game Master Paradigm. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 196–201. Springer, Heidelberg (2007)
3. Tychsen, A., Hitchens, M., Brolund, T., Kavakli, M.: The Game Master. In: Australasian Conference on Interactive Entertainment, Sydney, Australia, pp. 215–222. Creativity & Cognition Studios Press (2005)
4. Tychsen, A., Smith, J.H., Hitchens, M., Tosca, S.: Communication in Multi-player Role Playing Games – The Effect of Medium. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 277–288. Springer, Heidelberg (2006)
5. Young, R.M., Riedl, M.O., Branly, M., Jhala, A., Martin, R.J., Saretto, C.J.: An Architecture for Integrating Plan-based Behavior Generation with Interactive Game Environments. *Journal of Game Development* 1, 51–70 (2004)
6. Cutumisu, M., Onuczko, C., McNaughton, M., Roy, T., Schaeffer, J., Schumacher, A., Siegel, J., Szafron, D., Waugh, K., Carbonaro, M., Duff, H., Gillis, S.: ScriptEase: A Generative/Adaptive Programming Paradigm for Game Scripting. *Science of Computer Programming* 67(1)
7. Peinado, F., Navarro, Á., Gervás, P.: A Testbed Environment for Interactive Storytellers. In: International Conference on Intelligent Technologies for Interactive Entertainment, ACM Digital Library, New York (2008)
8. JBoss: Drools (2008), <http://www.jboss.org/drools/>

StoryTec: A Digital Storytelling Platform for the Authoring and Experiencing of Interactive and Non-linear Stories

Stefan Göbel, Luca Salvatore, Robert Arthur Konrad, and Florian Mehm

ZGDV Darmstadt e.V., Digital Storytelling Group,
64283 Darmstadt, Germany
{stefan.goebel, luca.salvatore, rakonrad, fmehm}@zgdv.de
<http://www.zgdv.de/distel/>

Abstract. Based on the results of INSCAPE and U-CREATE, StoryTec is introduced as a Digital Storytelling platform for the authoring and experiencing of interactive, non-linear stories. Apart from a description of StoryTec this paper focuses on the platform's authoring and runtime environment.

1 Introduction

Recently, within the information society and the up-coming, prospering creative industries, a lot of research has been investigated into Storytelling systems and authoring tools for the creation and representation of Story based scenarios in a broad range of application domains, for instance marketing and advertisement, training and simulation, E-Learning, Edutainment [6], educational games [7] or any kind of interactive digital media appliances. Prominent examples are the EU funded projects U-CREATE¹, nm2² and INSCAPE³, whose global aim was to allow even non-specialists the easy and efficient creation of interactive stories for various application domains.

With respect to content production and the authoring process, a lot of existing tools are available on the market, e.g. Flash, Director, 3DS, Maya, Toonz, Virtools, Quest3D, Blender or GameStudio, which are commonly used to create interactive stories respectively multimedia appliances and digital games. Most of them are restricted to a linear story representation and there is a lack of instruments to organize and structure stories or to define object behaviors and interactivity in general. Hence, authoring complex, interactive, non-linear stories with many story branches and interactive story units might become -especially for people without programming skills- a long, challenging and confusing process. For that, projects such as U-CREATE [4] and INSCAPE [1, 5] offer an interactive visual representation of the story as part of a user-friendly and intuitive GUI [1, 2].

¹ U-CREATE: Authoring Tools for Edutainment Applications <http://www.u-create.org>

² nm2: new millenium, new media (FP6, IP), <http://www.ist-nm2.org/index.html>

³ INSCAPE: Interactive Storytelling for Creative People <http://www.inscapers.com>

2 StoryTec Platform

The StoryTec platform has been conceptualized and implemented by the Digital Storytelling group at ZGDV Darmstadt based on its results and lessons learned within INSCAPE, U-CREATE and several other projects. The StoryTec platform consists of two major components: An authoring environment and a runtime engine.

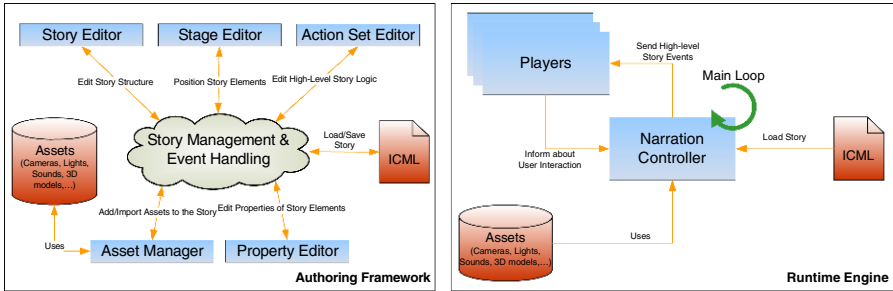


Fig. 1. Overview of the authoring and runtime parts of the StoryTec Platform

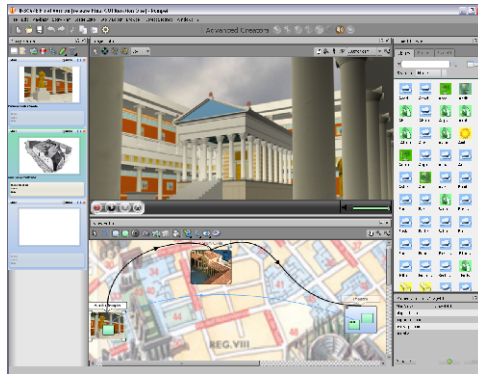


Fig. 2. The INSCAPE GUI

2.1 Authoring Platform

The authoring part is based on a pluggable framework and is composed of a Story Editor, a Stage Editor, an Action Set Editor, a Property Editor and an Asset Manager.

The main purpose of the framework is to connect these components and provide a common GUI. Furthermore, the framework is responsible for managing and keeping the story structure up-to-date.

Story Editor: The Story Editor is used to manage the story structure. A story is structured by a hierarchically organized graph consisting of *scenes* and *complex scenes*. A complex scene typically defines a particular location within the story world and contains the physical surroundings (decor, non-interactive elements) which make up the background of the story. The Story Editor provides a representation of the story

graph, containing the scenes/complex scenes and the transitions between these story units in a user-defined arrangement. This supports the user to recognize the story as a whole and may provide hints about the main features and the author's intention. A semantic zoom feature with different (level-of-detail) representations for the story elements helps to prevent the representation from getting cluttered. Usability studies carried out within the INSCAPE project pointed out the potential of the Story Editor, highlighting it as an innovative feature not present in comparable software [1, 2].

Stage Editor: Similar to video game level editors such as *Neverwinter Nights*, the Stage Editor represents the place to create scenes by inserting objects from the Asset Manager (objects library, using drag-and-drop). All objects with a physical representation (props, etc.) are visible and can be manipulated. Since the authoring environment is based on a pluggable framework, it is possible to have different Stage Editors for different types of application (3D or 2D). A Nebula⁴-based 3D Stage Editor has been implemented, further interfaces to GameStudio⁵, Blender⁶ and Flash are currently under development.

Action Set Editor: The Action Set Editor provides a visual environment for defining high-level story logic for every scene. The logic is described by a set of rules composed of actions (e.g. walkTo 'Door') and conditions (e.g. 'if player1 enters the castle'). To allow non-programmers to define the story logic, an interactive 2D graph visualization based on the UML Activity Diagram⁷ is provided in contrast to a script-based approach in traditional authoring tools (and game editors) or the INSCAPE framework.

Asset Manager: The Asset Manager is responsible for importing various types of assets (cameras, lights, props, etc.) to its library. Assets are visualized by icons and preview images in a separate frame. The Asset Manager enables dragging of assets into a scene or complex scene in the Story Editor or into the Stage Editor

2.2 Runtime Engine

The runtime part of the StoryTec platform is responsible for executing the stories. A Narration Controller (NC) forms the core of the runtime engine. It interprets and executes the created stories encoded in the XML-based format ICML conceptualised by the Storytelling group at ZGDV during the INSCAPE project [4]. The graph is traversed by following the transitions from scene to scene. If a transition is executed the NC informs the Player component about a scene change. On the other side, the Player informs the NC about user events (e.g. 'Position of prop_1 changed'). Taking into account these events, the NC traverses every scene's action set, evaluates the conditions and informs the Player about the actions to be executed (e.g. 'Open car door'). Hereby, a strategies concept being elaborated within the INSCAPEs' *Story Pacing* research topic [5] provides the basis for adaptive story control.

⁴ Nebula game engine, open source, <http://www.radonlabs.de/technologynebula2.html>

⁵ GameStudio, <http://www.gamestudio.de/>

⁶ Blender is an open source 3D content creation suite with a Game Engine (www.blender.org).

⁷ UML Activity Diagram is part of the UML 2.0 specification (www.uml.org).

3 Conclusion

This demonstration paper describes the Storytelling platform StoryTec, consisting of a comprehensive authoring framework with editors (Story Editor, Stage Editor, Action Set Editor, etc.) enabling authors without programming skills to create interactive stories and a runtime engine responsible for fluent story control. An extended, comprehensive description of the StoryTec platform and underlying concepts provide Göbel et al. in [3].

Within 80Days⁸, the StoryTec platform will be further cultivated and adapted for the creation and control of interactive, Story-based educational games.

References

1. Balet, O.: INSCAPE An Authoring Platform for Interactive Storytelling. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 176–177. Springer, Heidelberg (2007)
2. Dade-Robertson, M.: Visual Scenario Representation in the Context of a Tool for Interactive Storytelling. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 3–12. Springer, Heidelberg (2007)
3. Göbel, S., Salvatore, L., Konrad, R.: StoryTec: A Digital Storytelling Platform for the Authoring and Experiencing of Interactive and Non-linear Stories. In: AXMEDIS 2008, 4th International Conference on Automated Solutions for Cross Media Content and Multi-channel Distribution (2008)
4. Göbel, S., Salvatore, L., Konrad, R., Sauer, S., Osswald, K.: U-CREATE: Authoring Tool for the Creation of Interactive Storytelling-based Edutainment Applications. In: Cappellini, V., Hemsley, J. (eds.) EVA 2007 Florence Proceedings, pp. 53–58. Pitagora Editrice Bologna, Bologna (2007)
5. Göbel, S., Malkewitz, R., Becker, F.: Story Pacing in Interactive Storytelling. In: Pan, Z., Aylett, R.S., Diener, H., Jin, X., Göbel, S., Li, L. (eds.) Edutainment 2006. LNCS, vol. 3942, pp. 419–428. Springer, Heidelberg (2006)
6. Hoffmann, A., Göbel, S., Schneider, O., Iurgel, I.: Storytelling based Edutainment Applications. In: Tan, L., Subramaniam, R. (eds.) E-Learning and Virtual Science Centers, pp. 190–214. Information Science Publishing, Hershey (2005)
7. Kickmeier-Rust, M.D., Peirce, N., Conlan, O., Schwarz, D., Verpoorten, D., Albert, D.: Immersive Digital Games: The Interfaces for Next-Generation E-Learning? In: Stephanidis, C. (ed.) HCI 2007. LNCS, vol. 4556, pp. 647–656. Springer, Heidelberg (2007)

⁸ 80Days: Around an inspiring virtual learning world in eighty days. FP7, STREP, IST-2007-215918, challenge 4.1.2 Technologies enhanced learning. www.eightydays.eu

Workshop: Impro Theatre

Dagmar Dörger and Martin Geisler

FH Erfurt, University of Applied Sciences, Altonaer Str. 25, 99085 Erfurt, Germany
doerger@fh-erfurt.de, martin.geisler@fh-erfurt.de

Workshop Abstract

Interactive storytelling is also possible without a computer. Even a theatre play doesn't need to be based on a predefined text in order to be performed. A storyline in an improvisational play evolves dynamically from the input of the audience, from actions of the facilitator and from spontaneous ideas of the actors and actresses. The audience becomes an interactive element and – potentially – the actual author of the story. However, this doesn't mean that an Impro theatre play would not need a predefined concept. Also, the players and performers need at least basic training in the skills for impromptu acting. Even the development of a story that unfolds right in the moment of the performance needs basic assumptions. [1][2]

This workshop introduces basic rules, strategies and concepts that are useful for improvisers. The covered topics include

- **rules and strategies of traditional Impro theatre**, such as preparations, concepts of how to solve problems on stage, and consequences for modern interactive storytelling, as well as
- **insights into the development of a storyline** through a concrete Impro theatre play performed by the Erfurt group 'ImproVision'.

The workshop consists of the following four parts:

- **Introduction:** Background information on Impro theatre is provided in general, and particularly on the technique of an "Autorenspiel" (author play).
- **Performance:** During the performance of an "Autorenspiel", the audience becomes author of the play/story, and workshop participants experience how their ideas are transformed into this unique play.
- **Basics and Techniques:** Based on the outcome of the performance, the resulting story is analysed, in terms of its development and underlying structures, the techniques and possibilities of interaction, as well as limitations and problems on stage.
- **Meta-Discussion:** The overlapping aspects of digital and non-digital interactive storytelling are discussed.

References

1. Dörger, D., Nickel, H.-W.: Improvisationstheater – ein Überblick. Das Publikum als Auto. Schibri Verlag, Berlin (2008) (in German)
2. Johnstone, K.: Impro for Storytellers. Theatresports and the Art of Making Things Happen. Faber and Faber, London (1999)

Workshop: Pen-and-Paper Role-Playing

Florian Berger and Alexander Marbach

Pädagogische Hochschule / University of Education Weingarten, Germany
{berger,marbach}@md-phw.de

Workshop Abstract

Over the last 30 years, a phenomenon of human-to-human interactive storytelling has developed, known as *fantasy role-playing games (RPG)*, beginning with the publishing of the game *Dungeons & Dragons* in 1974. Since then, a role-playing culture has emerged, incorporating concepts from theatre, board games, history reenactment and storytelling, remixed in an almost post-modern way.

A role-playing session is in its very nature an interactive storytelling session, and the *game master* usually solves problems similar to those involved in Interactive Digital Storytelling (IDS) systems.

The goal of this workshop is to offer researchers in IDS valuable insights into principles of pen-and-paper RPGs, in particular in strategies of the game master. This is why the authors bring such a session to the conference, and moreover involve workshop participants as active players. The workshop consists of the following parts:

Introduction: The basic concepts of a pen-and-paper RPG are explained: the *game master* as the creator and narrator of the virtual world; *players* as the leaders of their respective fictional characters; resolving conflicts using rulesets; rules of in-game communication; character creation.

Playing Session: This is the actual role-playing, with the players completing a quest created by the game master. He has prepared a set of *non-player characters (NPCs)*, maps of some *key locations* and some *key events*, but the actual storyline emerges by player interaction over the course of the game. Typical RPG features include *suspense*, *time pressure*, *riddles*, *teamwork*, *acting* and *conflicts/combat*. There are some typical problems a game master has to face during gameplay, most of them related to the conflict between the players' freedom to act and the demand for a consistent storyline. A game master makes use of several techniques to tackle these challenges.

Meta Discussion: This serves the purpose of answering questions for first-time players, but every participant is encouraged to give personal views and impressions and ask questions. Game mastering techniques like reacting towards unexpected developments are revealed and discussed. Possible implications of the game play for future research in interactive digital storytelling are summed up and discussed.

Workshop and Panel: The Authoring Process in Interactive Storytelling

Ulrike Spierling¹ and Ido Iurgel²

¹ FH Erfurt, University of Applied Sciences, Altonaerstr. 25, 99085 Erfurt, Germany

² CCG, Centro de Computação Gráfica and Centro Algoritmi, Universidade do Minho, Campus de Azurém, 4800 Guimarães, Portugal
spierling@fh-erfurt.de, ido.iurgel@ccg.pt

Workshop and Panel Abstract

At TIDSE 2006, we conducted a pre-conference demo workshop [1] that addressed the issue of developing accessible solutions for the creation of interactive storytelling artworks and applications. As a principle of the workshop, participants demonstrated their authoring systems using a shared story: the Grimm's fairy tale of Little Red Riding Hood. They were allowed to change and adapt the initial story in order to meet the underlying technical philosophy of their system, and in order to introduce interactivity. This exercise quickly made obvious that a) – not surprisingly – there are several ways and strategies to make a story interactive, and b) that the creative approaches suggested by these systems differ fundamentally, making it hard to get an integrated picture of 'Interactive Storytelling' from a creator's point of view.

The comparative approach of this 'hands-on' workshop prompted vivid discussions, and provided a platform to even explain some cryptic algorithms to a broader audience than only to A.I. researchers. Because of the encouraging results, the concept of the initial workshop is revived at Interactive Storytelling '08. In this second edition, in addition to tool demonstrations in a workshop, a plenary panel with brief presentations is organized. It provides a quick overview of showcases in current state-of-the-art authoring systems and creation methods.

The questions answered and topics discussed by each presenter of a tool are:

- **Tool Architecture:** Describe particular features of your approach or your technology for Interactive Storytelling!
- **Storyworld Concept:** Describe an outline of your particular adaptation of Little Red Riding Hood, including the role of the user / player!
- **Example Scene:** Describe in more detail one example scene of the conceived storyworld that is characteristic for your approach!
- **Creation Process:** Describe how an author uses your tools to create exactly that example scene!

The results are discussed further on in the form of a weblog [2].

References

1. Spierling, U., Iurgel, I.: Pre-Conference Demo Workshop "Little Red Cap": The Authoring Process in Interactive Storytelling. In: Göbel, S., Malkewitz, R., Iurgel, I. (eds.) TIDSE 2006. LNCS, vol. 4326, pp. 193–194. Springer, Heidelberg (2006)
2. Little Red Cap Authoring Tools Weblog,
<http://www.fh-erfurt.de/ai/icids08/redblog>

Author Index

- André, Elisabeth 14
Arinbjarnar, Maria 180
Aylett, Ruth 218, 273
- Bae, Byung-Chull 156
Barber, Heather 214
Bendon, Helen 41
Berger, Florian 92, 330
Berndt, Axel 126, 132
Bída, Michal 59
Bogost, Ian 268
Böse, Ralf 114
Boyd Davis, Stephen 41
Brisson, Antonio 218
Brolund, Thea 242
Brom, Cyril 59
Bulitko, Vadim 230
Burkert, Ondřej 59
- Camanho, Marcelo M. 198
Cavazza, Marc 14, 83, 285
Champagnat, Ronan 14
Cheong, Yun-Gyung 144, 297
Christie, Marc 14
Ciarlini, Angelo E.M. 198
Cortiguera, Héctor 321
- Donikian, Stéphane 14
Dörger, Dagmar 329
Dória, Thiago R. 198
- Feijó, Bruno 198
Figueiredo, Rui 218
Fischer, Ulrich 55
Friess, Regina 108
Furst, Merrick L. 268
Furtado, Antonio L. 198
- Geisler, Martin 329
Gemrot, Jakub 59
Göbel, Stefan 325
Gordon, Andrew S. 32
Grønbaek, Kaj 20
- Hansen, Frank Allan 20
Hartmann, Knut 126
- Hartmann, Tilo 14
Hitchens, Michael 242
Horttana, Tommi 44
- Isbell, Charles 268
Iurgel, Ido A. 89, 331
- Jain, Abhishek 186
Jhala, Arnav 210
- Kadlec, Rudolf 59
Kim, Jin-Young 297
Kim, Yeo-Jin 297
Klimmt, Christoph 14
Konrad, Robert Arthur 325
Kortbek, Karen Johanne 20
Kriegel, Michael 273
Kruizinga, Edze 264
Kudenko, Daniel 180, 214
- Lankoski, Petri 44
León, Carlos 321
Louchart, Sandy 273
- Marbach, Alexander 92, 330
Mehm, Florian 325
Mehta, Manish 186
Milam, David 96
Min, Wook-Hee 297
Moar, Magnus 41
Müller, Wolfgang 92
- Navarro, Álvaro 321
- Olivier, Patrick 14
Ontañón, Santiago 186
- Paiva, Ana 218
Parry, Nye 41
Peinado, Federico 83, 321
Petta, Paolo 14
Pinchbeck, Dan 51
Pizzi, David 83, 285
Pozzer, Cesar T. 198
- Ram, Ashwin 186
Riedl, Mark O. 168, 268

- Roberts, David L. 268
Ryan, Marie-Laure 6
Salvatore, Luca 325
Seif El-Nasr, Magy 71, 96
Shim, Eok-Soo 297
Spetch, Marcia 230
Spierling, Ulrike 14, 114, 331
Stern, Andrew 1
Struck, Georg 114
Sugandh, Neha 168
Sumi, Kaoru 48
Swanson, Reid 32
Swartjes, Ivo 264, 273
Szilas, Nicolas 14
Tanenbaum, Theresa Jean 250
Tanenbaum, Karen 250
Theisel, Holger 132
Theune, Mariët 264
Thue, David 230
Tychsen, Anders 242, 309
Vorderer, Peter 14
Wakkary, Ron 96
Wei, Huaxin 71
Weiß, Sebastian A. 92
Young, R. Michael 144, 156