Chapter 10
Emotion-Driven Narrative Generation

Brian O'Neill and Mark Riedl

Abstract While a number of systems have been developed that can generate stories, the challenge of generating stories that elicit emotions from human audiences remains an open problem. With the development of models of emotion, it would be possible to use these models as means of evaluating stories for their emotional content. In this chapter, we discuss Dramatis, a model of suspense. This model measures the level of suspense in a story by attempting to determine the best method for the protagonist to avoid a negative outcome. We discuss the possibilities for Dramatis and other emotion models for improving intelligent generation of narratives.

Introduction

Games are one of several common forms of entertainment that makes use of narrative. Many game genres use fictional context to reinforce the immersion within the game world and to motivate the player's activities. These fictional contexts answer the question, "Why am I, the player, engaging in a particular activity?" The fictional context may further induce an affective response from a player: dramatic tension over how events are unfolding, strong positive or negative feelings towards virtual characters, or suspense over what might happen next.

In many game narratives, the narrative was pre-determined by the authors and designers. The player has little or no capacity for affecting the events of the story, because the game has a single, linear narrative arc. In contrast, interactive narrative is a form of digital interactive experience in which users create or influence a dramatic storyline through actions, either by assuming the role of a character in

B. O'Neill (✉)
Department of Computer Science and Information Technology, Western New England University, 1215 Wilbraham Road, Springfield, MA 01119, USA
e-mail: brian.oneill@wne.edu

M. Riedl
School of Interactive Computing, Georgia Institute of Technology, 85 Fifth Street, Atlanta, GA 30308, USA
e-mail: riedl@cc.gatech.edu

K. Karpouzis, G.N. Yannakakis (eds.), *Emotion in Games*, Socio-Affective Computing 4, DOI 10.1007/978-3-319-41316-7_10

a fictional virtual world, issuing commands to computer-controlled characters, or directly manipulating the fictional world state. The simplest form of an interactive narrative is a branching story, such as Choose-Your-Own-Adventure books and hypertexts, in which plot points are followed by a number of options that lead to different, alternative narratives unfolding. More complex interactive narrative systems use artificial intelligence (AI) to construct the story on the fly in accordance with the player's desires. In AI-driven interactive narrative, a *drama manager*—an omniscient background agent that monitors the state of the fictional world—conducts a search through possible future narrative trajectories, determines what will happen next in the game, and coordinates virtual characters to bring about the best narrative possible. The difficulty lies within recognizing which of the trajectories is most interesting to the player.

Good narratives are not simply a series of events—good narratives elicit emotional responses from their audiences. However, generating stories that intentionally elicit an emotional response is challenging. Maintaining an emotional level in an interactive narrative requires interference from an intelligent manager, keeping the story on trajectories that are expected to keep emotional content high. Models of emotion could help address these issues. With an understanding of emotional responses to stories, these models could be developed and used to generate emotion-inducing stories. In this chapter, we describe Dramatis, a model of suspense. With such a model, we can judge the level of suspense a reader or player would feel from following the narrative of a particular trajectory. We will also discuss the consequences for having such a model (or models of other emotions) for story generation and interactive narrative.

Background

Before discussing the Dramatis system or how we can judge the suspense level of a particular future trajectory, we must clarify what we mean by narrative and suspense. We will also briefly discuss story generation and interactive narrative, while providing examples of such systems that have attempted to address suspense and related emotional responses.

Narrative

Narrative is ubiquitous in human culture. Narratives are used in a variety of forms of entertainment, including books, films, and games. In addition to entertainment, people create and share narratives in order to explain the world around them. Prince defines narrative as follows [25]:

Narrative: The representation ...of one or more real or fictive events communicated by one, two, or several (more or less overt) narrators to one, two, or several (more or less overt) naratees.

Put another way, a narrative is the communication of events (rather than simple facts) from a narrator to a reader or listener. The key to the definition is the requirement for events. A fact (e.g., "It is snowing.") does not constitute a narrative. By contrast, a single event ("I went to the store.") is a narrative, albeit not one that is particularly interesting. The "main incidents" of a narrative come together to form a plot. Plots may follow a common structure, such as the traditional Aristotelian arc [2], or Freytag's triangle [11].

Narratologists distinguish between the events of the story (known as the *fabula*) and the presentation of those events by the narrator(s) to the narratee(s) (known as the *sjuzhet*). The *fabula* contains all events of the narrative, regardless of the order of presentation, or whether they are presented to the audience at all. A *sjuzhet* is a particular ordering and presentation of a subset of the events contained in the fabula. It is common for the sjuzhet to exclude events from the fabula, or to alter the order of events. A fabula may therefore have multiple sjuzhets, depending on what the narrator chooses to include from the set of events and the order in which they are told. As an example of fabula and sjuzhet, consider the original *Back to the Future* film. The events of the film, including those not shown on-screen, make up the fabula. The presentation shown to the audience is one sjuzhet. An alternate sjuzhet would show the events in chronological order.

Story Generation and Interactive Narrative

Story generation refers to the ability of artificial intelligence to create new stories. Computational approaches to story generation largely take one of two approaches: search-based approaches [17, 23, 27, 31] or adaptive approaches [13, 22, 30]. Search-based systems explore a space of possible sequences of actions, using some heuristic of quality to compare them. Adaptive systems start from a library of known stories. These stories are modified or recombined into new stories, sometimes using analogical reasoning. In some cases, story generation systems work to create both a new fabula or sjuzhet. However, some systems focus on only creating one or the other (e.g. generating a sjuzhet for a given fabula).

Interactive narrative is a form of story generation that features the audience as a user who can influence the narrative as it progresses [26]. The user, typically in the role of the protagonist, can take actions in the story-world, thereby affecting the path and outcome of the story. Some interactive narratives give the user control over the world as an observer, rather than giving them direct control of a character. Non-player characters (NPCs) in the interactive narrative may be controlled by an experience manager, which affects the world in order to maintain story quality. The definition of quality varies depending on the particular system, but may include closeness to the author's intended story or measures of player experience, such as the expected emotional impact on the player.

A number of story generation and interactive narrative systems have attempted to address audience emotion in some way. Suspenser [7] is a story generation system which, given a fabula, attempts to identify the most suspenseful sjuzhet. Similarly, Prevoyant [3] uses flashbacks and foreshadowing to reorder a story with the goal of creating the most surprising sjuzhet from a fabula. Ware et al. [31] developed a model of narrative conflict using planning, applicable to story generation and interactive narrative. Façade [16] and Merchant of Venice [24] are interactive narrative systems which establish ideal tension curves for the narrative. The drama manager in each system affects the story by trying to get the tension in the story to the pre-defined level. Other story generation and interactive narrative systems [4, 22] also use tension as a metric, though there is no consistent definition of tension among them.

Suspense

Expert storytellers who craft their narratives for entertainment often structure their sjuzhets with the intent of eliciting emotional responses from the narratees (readers, game players, film viewers, etc.). The idea that story structure is correlated with audience enjoyment dates to Aristotle [2]. Suspense is one of many commonly used tools for creating emotional responses and has been found to contribute to reader enjoyment [29]. There are many definitions of suspense, coming from the fields of narratology [1, 6, 29], psychology [8, 12, 21], and entertainment theory [33], to name a few. Rather than consider each of these definitions, we will highlight the similarities in those definitions.

There are four attributes that are common among the various definitions of suspense: (1) uncertainty about an outcome, (2) a particularly desirable or undesirable possible result to that uncertain outcome, (3) an audience affinity for the character whose outcome is uncertain, and (4) a disparity of knowledge between the characters and the audience. The uncertainty of an outcome is the most important feature, and there must be meaning behind this uncertainty. There must be a substantial possibility of an undesirable state resulting for the character. However, there is no suspense unless the audience cares about the character. If the audience does not like the character or cannot identify with the character, then they will not feel suspense about the character's outcome. Finally, suspense can be generated by giving the audience more knowledge about a situation than the characters have. In such cases, the audience will be aware of the potential dire consequences of a situation, while the characters may have no idea.

One definition of suspense that we wish to highlight comes from Gerrig and Bernardo [12]

Readers feel suspense when led to believe that the quantity or quality of paths through the hero's problem space has become diminished.

Gerrig and Bernardo generated this definition by studying the levels of suspense self-reported by readers given different versions of story excerpts. Readers act as problem-solvers on behalf of the protagonist, attempting to identify solutions that avert a negative outcome for the protagonist. When readers struggle to find solutions, or only find low-quality solutions, readers perceive more suspense. Thus, in a sense, readers find themselves in an interactive narrative in their own mind, as they evaluate potential future trajectories for the protagonist. However, unlike in an interactive narrative, the reader lacks the ability to decide for the protagonist which path to take.

Dramatis adopts a reformulation of Gerrig and Bernardo's definition. This reformulation is discussed in the "Reformulating Gerrig and Bernardo" section below. Suspenser [7] applies Gerrig and Bernardo's definition of suspense in its attempt to find the most suspenseful sjuzhet for a given fabula. The system measures suspense by projecting all possible future plans and determining the ratio of failed plans to successful plans, with suspense increasing as the ratio increases.

Dramatis

Dramatis is a computational model of suspense felt by the reader of a story. The model reads a story and calculates the level of suspense over time. The model uses a reformulation of Gerrig and Bernardo's definition of suspense (details of which are discussed in the next section). Dramatis reads a discretized symbolic-logic version of story events, determines whether characters are facing an undesirable outcome, and generates and evaluates the quality of the best plan for avoiding that outcome. The evaluation of quality is correlated with the level of suspense at that moment of the story.

Reformulating Gerrig and Bernardo

Recall Gerrig and Bernardo's definition of suspense, introduced above: "Readers feel suspense when led to believe that the quantity or quality of paths through the hero's problem space has become diminished." Gerrig and Bernardo describe a search space, where the search is conducted by the reader on behalf of the hero of the story. The search space consists of possible future states of the story world. Readers, therefore, are searching through a series of potential storylines and judging which one is best. Suspense is generated, in part, by how authors manipulate the space.

How do authors induce suspense? According to Gerrig and Bernardo's definition and their studies of readers, authors can manipulate the quantity and/or quality of paths in the hero's space. Authors may propose possible solutions, potentially implying an increase in quantity or that the plan is high quality, before striking it down, thereby diminishing the quantity of plans available. Authors may otherwise indicate suggest plans that they know to be faulty in order to distract readers from the solution that will ultimately be used by the hero. So authors create suspense by manipulating the search space or how the reader traverses the search space.

Gerrig and Bernardo's definition of suspense is computationally intractable. It is not possible to measure all paths through the problem space in terms of success and failure and weigh the ratio, because paths may fail as a result of the search process or the planning problem, rather than as a result of the conditions of the story world. Additionally, Gerrig and Bernardo's definition suggests that humans regenerate the search space repeatedly while reading. However, many of the definitions of suspense indicated that humans only search the space when prompted to do so by a potential undesirable outcome. Further, regenerating the search space requires the ability to identify the causal consequences of story events, an inference that can only occur when the reader puts the story aside [14]. Finally, human memory is resource-bounded, and they are therefore incapable of considering the entire space of possible events, let alone constantly regenerating that search space as they read.

As a consequence, we reformulate Gerrig and Bernardo's definition of suspense as follows:

> Given the belief that a character can face a negative outcome, one can assume that this outcome will occur and search for the single most likely plan—the *escape plan*—in which the protagonist avoids this outcome.

Gerrig and Bernardo refer to the quality of paths through the hero's problem space, though they are not precise with how quality is measured. We consider the escape plan's quality to be its perceived likelihood of success from the perspective of the reader. Using perceived likelihood (rather than actual likelihood) allows us to account for ways in which the author might manipulate the problem space, as well as account for the disparity in knowledge between the characters and the audience. We use a model of reader memory to calculate the perceived likelihood, working from the concept that humans consider the first thought retrieved from memory to be the most likely thing to actually occur [15]. Additionally, this reformulation requires neither constant regeneration of the search space, nor generation of the total search space. Finally, by searching for a single escape plan, there is no comparison between the number of successful plans and the number of failed plans.

Dramatis Algorithm and Inputs

Figure 10.1 shows the Dramatis algorithm for measuring suspense in a story. Subsequent sections break down these steps in further detail. Dramatis reads stories in a discretized symbolic format, which we call *time-slices*. Each time-slice describes one action in the story and provides state information about the characters and location of the scene. As Dramatis reads the story, it searches a library of scripts to identify one whose sequence of events matches the events observed in the story-so-far. The script provides information about what negative outcomes may occur in the near future. Once a negative outcome has been predicted, Dramatis generates an escape plan to avert the negative outcome. The perceived likelihood of the escape plan's success is correlated with the level of suspense at that moment. As the story continues and new information is gained in the story, Dramatis revises its escape plans and its measures of likelihood and suspense, potentially generating new plans as old escape plans cease to be viable. At the conclusion of the story, we are left with a curve showing suspense over time.

Let SL be a script library, OL be an operator library, and T be an ordered sequence of time-slices. Let NM represent narrative memory, MS be a model of reader salience, and EP be an escape plan.

DRAMATIS(SL, OL, T):
 $NM \leftarrow \emptyset$
 $MS \leftarrow \emptyset$
 $EP \leftarrow \emptyset$
 $EscapePlans[] \leftarrow \emptyset$
 $Suspense[] \leftarrow \emptyset$
 For each time-slice $t \in T$:
  Read t into NM
  $Scr \leftarrow$ Retrieve-Script(NM, SL)
  $MS \leftarrow$ Update-Salience-Model(NM, Scr, MS)
  If t is the next action in EP:
   $cost \leftarrow$ Recalculate-Plan-Cost(EP, MS)
  Else:
   $Links[] \leftarrow$ Identify-Links-To-Break(Scr, NM)
   For each link $L \in Links$:
    $\langle plans[L], costs[L] \rangle \leftarrow$ Generate-Escape-Plan(L, MS, NM, OL)
   $cost \leftarrow \min(costs[L])$
   Let EP be the plan in $plans[]$ with minimum cost in $costs[]$
  $EscapePlans[t] \leftarrow EP$
  $Suspense[t] \leftarrow cost$
 Return $\langle EscapePlans, Suspense \rangle$

Fig. 10.1 Dramatis algorithm

```
Action: deliver-food (Waitress, vodkaMartini, James_Bond)
Location: Casino Royale
Characters: Waitress, James_Bond
Effects: at(James_Bond pokerTable) place(pokerTable)
```

Fig. 10.2 Example time-slice

Time-Slices

As input, Dramatis requires a story, which it receives as an ordered set of discretized time-slices. Each time-slice describes exactly one action in the story. Time-Slices contain a representation of the event in the form of an instantiated STRIPS operator, the characters in the scene, the location of the scene, and the effects of the action that occurred. After the time-slice is read, its contents are stored in narrative memory, where Dramatis tracks the state of the story-world as it continues reading. In some cases, a time-slice may also contain a reference to an opposing character's plan. These plans are used to infer possible negative outcomes that the protagonist may face.

Figure 10.2 shows an example time-slice. In this time-slice, the STRIPS operator is named deliver-food, with the parameters `Waitress`, `vodkaMartini`, and `James_Bond`. The location of the time-slice is `Casino Royale`, while `Waitress` and `James_Bond` are annotated as characters. The effects listed will be added to narrative memory, along with the expected effects of the deliver-food operator. This time-slice does not contain a reference to an opposing character's plan.

Scripts

In addition to the story, Dramatis receives a script library as input. In general, scripts are conceptual frameworks used by people to navigate everyday situations (e.g., ordering food from a restaurant) [28]. When we go to restaurants, we are familiar with the typical pattern of being seated, ordering drinks, ordering food, eating the food, getting the check, and paying the check. However, when someone tells a story about eating at a restaurant and leaves out of those steps from their story, we are capable of inferring that it occurred. For example, when hearing a story about someone eating steak at a restaurant, we are capable of inferring that the steak had previously been ordered. This capacity for inference comes from scripts. Scripts, therefore, are inherently useful in understanding stories [9].

Scripts are typically learned through personal experience. However, people may also learn scripts through second-hand experiences or by hearing stories. For example, one might learn, through reading several fairy tales, that princes rescue damsels-in-distress. Thus, one could develop a script in which if a damsel is in distress, then she may be saved by a prince.

Traditionally, scripts are linear structures [28]. The scripts used by Dramatis are more akin to graphs, containing multiple possible paths that a story could take. Nodes in the graph represent events, while edges may represent temporal links or causal links. A temporal link is a directed edge in the graph, indicating that the source node occurs before the destination node in the graph. A causal link, also a directed edge, indicates that the source node provides one of the causal conditions necessary for the event in the destination node. The link is further annotated with what that condition is. Dramatis will use these scripts to infer negative outcomes (see "Predicting Negative Outcomes"). Nodes containing negative outcomes are annotated as such. Additionally, the scripts will be used in order to determine the goal situation for generating an escape plan (see "Generating Escape Plans").

Planning Operators

Finally, Dramatis is given a set of STRIPS planning operators [10] as input. STRIPS operators are described by an action, as well as the actors and objects needed to complete them. The operators define what conditions must be true in the world before the action can occur and what new conditions are true once the action has been completed. We use STRIPS operators for two purposes. First, each node in the script library represents an action and can therefore be bound to a STRIPS operator. Second, we use STRIPS operators to represent the actions that are available to characters when developing an escape plan.

Predicting Negative Outcomes

Dramatis reads a given story one time-slice at a time. As Dramatis reads, it attempts to predict whether the reader should expect a negative outcome for the protagonist. To make this prediction, Dramatis attempts to match the observed sequence of events (collected from the time-slices) to one of the scripts in its library. When choosing between scripts, Dramatis prefers scripts that make use of recently-observed story events and contains actions matching those observed earlier in the story. As the story continues, Dramatis tracks the script, maintaining a reference to the event in the script that was most recently observed in the story. Dramatis identifies the potential failure by traversing the script graph until a node is reached that is labeled as a negative outcome. Dramatis can conduct a similar process that makes use of knowledge of the antagonist's plans (when such plans have been relayed to the audience) rather than using a script. In such instances, the antagonist's plan is treated as though it were a script.

Measuring Reader Salience

After reading a time-slice, Dramatis adds each story element in the time-slice (characters, objects, locations, and events) to a model of reader memory. Dramatis uses the *Modified Event Indexing with Prediction* (MEI-P) model, which is based on previous psychological theories of the mental models of readers. Zwaan et al. [34] developed the Event Indexing (EI) model in order to understand how readers' conceptualizations of a story changed while they read. As part of his INFER system, Niehaus created the Modified Event Indexing (MEI) model [18] to account for narrative focus and readers' ability to draw inferences about the story while reading. MEI-P extends the MEI model by representing possible future events which are extracted from scripts.

MEI-P is a spreading activation network, where greater activation of a story element implies that the element is more salient in reader memory. Thus, a story element that has greater activation is more easily retrieved from memory. This ease of retrieval will be a factor in calculating the cost of actions when generating escape plans.

After Dramatis reads a time-slice, it creates a new node in the network which represents the new event. This node is connected to nodes representing the other story elements in the time-slice, creating new nodes if necessary. Each new edge in the network is given a weight of 1.0, while older connections in the network see their weights decay over time. Node activations are recalculated after each time-slice by giving them an initial activation of 1.0 and iteratively spreading node weights according to edge weights until activation levels throughout the network stabilize.

MEI-P includes predicted future events as well as observed events. Predicted events come from the script containing the negative outcome. Any event in the script that may follow from the most recently observed event is included in the MEI-P network. The salience of a predicted event decreases with how far in the future they are expected to be from the most recent of the story, just as older events decay with distance from the current event.

Generating Escape Plans

Let us review why Dramatis generates escape plans. The reformulation of Gerrig and Bernardo's definition of suspense states that when we can anticipate a negative outcome for the protagonist, the level of suspense will be correlated with the likelihood of the best plan that the reader (or Dramatis, simulating the role of the reader) can generate to avert the negative outcome. We determine the likelihood of a plan by determining the salience of story elements involved in the plan.

In order to generate an escape plan, we must first define the planning problem. The initial state of the planning problem is defined as the current state of the story world, constructed according to the information conveyed in time-slices. The goal situation is determined by identifying the causal links in the active script between the most recently observed event and the negative outcome. If any of those links were cut, then a causally necessary condition for the negative outcome would no longer be true in the story world. As a consequence, it would no longer be possible for the negative outcome to occur, at least on the current path through the script. Thus, the goal situation is determined by negating the causal conditions in the script. Thus, an escape plan is any series of actions that leads to a violation of one or more necessary conditions for the negative outcome, thereby averting that outcome.

Each STRIPS operator's cost is calculated using the level of activation of the corresponding nodes in the MEI-P salience model. This includes the activation of the characters and objects used in the operator, any locations referenced by the operator, the preconditions and effects of the event represented by the operator, and the activation of the event itself if it is part of the model. Recall that we assume that the plan that is most easily retrieved from memory will be perceived as the most likely plan to succeed [15]. Therefore, the plan that uses elements that are most easily retrieved will be considered most likely to succeed. As a result, the cost of an operator is inversely related to the activation of the story elements used by the operator. The cost of the plan is equal to the sum of the action costs. The level of suspense is equal to the total cost of the plan. Dramatis uses the Heuristic Search Planner (HSP) [5] to generate escape plans, though any planner that can return a near-optimal result for operators with non-uniform costs would suffice.

As Dramatis continues to read the story, it tracks the most recently generated escape plan. When newly observed events conform to the events predicted by the escape plan, then we recalculate the suspense level based on the remainder of the plan and the updated MEI-P salience model. Otherwise, Dramatis generates a new escape plan using the procedure described above. It is possible that Dramatis generates the same escape plan, even though the newest event was not part of the predicted escape plan. Thus, Dramatis generates an escape plan, and therefore a suspense rating, after each time-slice of the story. We use these suspense ratings to generate a *suspense curve*, showing the change in suspense over time. Dramatis was evaluated by comparing the generated suspense curves to ratings produced by humans reading text versions of the same stories [19, 20]. Figure 10.3 shows a suspense curve created by Dramatis during this evaluation.

Future of Emotion-Driven Story Generation

Emotional models, such as Dramatis, provide an opportunity to increase the emotional content of artificially generated stories and interactive narratives. Why should these narratives be emotion-driven? Emotional content is more interesting and more entertaining [29]. Without emotion, we have a sequence of boring events. Emotional content, including suspense, makes stories worth hearing, reading, and playing. Dramatis provides a model of suspense. With other models of emotion, it will be possible to intelligently author narratives that can induce the entire spectrum of human emotions from their audiences. Dramatis, and other models of

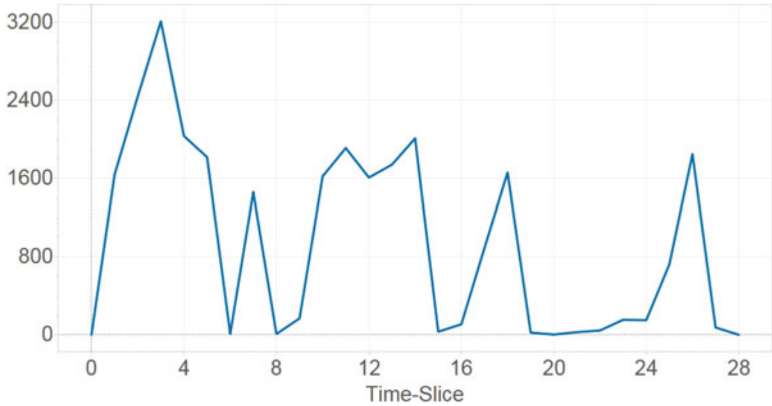

Fig. 10.3 A sample suspense curve created by Dramatis

emotion, constitute a means of direct, theory-driven evaluation of the quality of story content, under the Experience-Driven Procedural Content Generation (EDPCG) framework [32]. Under that same framework, Dramatis is an example of a model-based player—or perhaps in this case, audience—experience model. Dramatis does not address other aspects of the framework—most notably, content generation. However, Dramatis, or any other narrative emotion model, can easily serve as heuristic or evaluation function for a story generation or interactive narrative system.

Additionally, a story generation or interactive narrative system could leverage Dramatis by reasoning about the escape plans generated by the model. Suppose Dramatis were given an incomplete story to evaluate. It would read the story, generate an escape plan for the last known event of the story, and generate a suspense rating. The story generation system would then consider what events to add to the story-in-progress that would increase or decrease the level of suspense. Suspense is increased by inserting events that reduce the viability and perceived likelihood of the escape plan, while decreasing suspense requires making the escape plan seem more likely to succeed. In the case of story generation, it may be possible to insert or remove events from the middle of the incomplete story to alter the suspense level. (This would be less valuable for interactive narrative, as it would require going back in the story after the player has already progressed.) This process could continue iteratively, adding events and regenerating escape plans, until a minimal threshold of suspense is crossed, or until a pre-defined ideal suspense curve is matched. Future emotional models may have elements (akin to Dramatis escape plans) that allow for similar iterative processes.

Concluding Remarks

Stories that elicit emotion from their audiences are better, more entertaining stories. The ability to generate emotional stories automatically leads to the capacity for good stories-on-demand. This ability bodes well for games. When games can be created with emotional stories, both human-authored and intelligently generated, then it will be possible for games to be tailored to an individual player's emotions. A game can be made with an unlimited supply of stories, where each story affects the player's emotions differently, and without the need to author the stories well in advance.

References

1. Abbott HP (2008) The Cambridge introduction to narrative. Cambridge University Press, Cambridge
2. Aristotle (1992) The poetics. Prometheus Books, Buffalo
3. Bae B, Young RM (2008) A use of flashback and foreshadowing for surprise arousal in narrative using a plan-based approach. In: Proceedings of the 2008 international conference on interactive digital storytelling. Springer, Heidelberg, pp 156–167
4. Barber H, Kudenko D (2008) Generation of dilemma-based interactive narratives with a changeable story goal. In: Proceedings of the 2nd international conference on intelligent technologies for interactive entertainment. ICST, Brussels, pp 1–10
5. Bonet B, Geffner H (2001) Planning as heuristic search. Artif Intell 129:5–33
6. Branigan E (1992) Narrative comprehension and film. Routledge, New York
7. Cheong Y (2007) A computational model of narrative generation for suspense. PhD dissertation, North Carolina State University, Raleigh
8. Comisky P, Bryant J (1982) Factors involved in generating suspense. Hum Commun Res 9:49–58
9. Cullingford RE (1981) SAM and micro SAM. In: Schank R, Riesbeck CK (eds) Inside computer understanding. Lawrence Erlbaum Associates, Hillsdale
10. Fikes RE, Nilsson NJ (1971) STRIPS: a new approach to the application of theorem proving to problem solving. Artif Intell 2:189–208
11. Freytag G (1968) The technique of the drama: an exposition of dramatic composition and art. Johnston Reprint Corporation
12. Gerrig RJ, Bernardo ABI (1994) Readers as problem-solvers in the experience of suspense. Poetics 22:459–472
13. Gervás P, Díaz-Agudo B, Peinado F, Hervás R (2005) Story plot generation based on CBR. Knowl-Based Syst 18:235–242
14. Graesser AC, Singer M, Trabasso T (1994) Constructing inferences during narrative text comprehension. Psychol Rev 101:371–395
15. MacLeod C, Campbell L (1992) Memory accessibility and probability judgements: an experimental evaluation of the availability heuristic. J Personal Soc Psychol 63:890–902
16. Mateas M (2002) Interactive drama, art and artificial intelligence. PhD dissertation, Carnegie Mellon University, Pittsburgh
17. Meehan J (1981) TALE-SPIN. In: Schank RC, Riesbeck CK (eds) Inside computer understanding. Lawrence Erlbaum Associates, Hillsdale
18. Niehaus J (2009) Cognitive models of discourse comprehension for narrative generation. PhD dissertation, North Carolina State University, Raleigh
19. O'Neill B (2013) A computational model of suspense for the augmentation of intelligent story generation. PhD dissertation, Georgia Institute of Technology, Atlanta
20. O'Neill B, Riedl M (2014) Dramatis: a computational model of suspense. In: Proceedings of the twenty-eighth AAAI conference on artificial intelligence (AAAI-14). AAAI, Menlo Park, pp 944–950
21. Ortony A, Clore GL, Collins A (1988) The cognitive structure of emotions. Cambridge University Press, Cambridge

22. Pérez y Pérez R, Sharples M (2001) MEXICA: a computational model of a cognitive account of creative writing. J Exp Theor Artif Intell 13:119–139
23. Porteous J, Cavazza M (2009) Controlling narrative generation with planning trajectories. In: Proceedings of the 2nd international conference on interactive digital storytelling. Springer, Heidelberg, pp 234–245
24. Porteous J, Teutenberg J, Pizzi D, Cavazza M (2011) Visual programming of plan dynamics using constraints and landmarks. In: Proceedings of the 21st international conference on automated planning and scheduling. AAAI, Menlo Park, pp 186–193
25. Prince G (2003) Dictionary of narratology. University of Nebraska Press, Lincoln
26. Riedl MO, Bulitko V (2013) Interactive narrative: an intelligent systems approach. AI Mag 34:67–77
27. Riedl MO, Young RM (2010) Narrative planning: balancing plot and character. J Artif Intell Res 39:217–268
28. Schank RC, Abelson RP (1977) Scripts, plans, goals, and understanding: an inquiry into human knowledge structures. Lawrence Erlbaum Associates, Hillsdale
29. Tan ES (1996) Emotion and the structure of narrative film: film as an emotion machine. Routledge, New York
30. Turner SR (1993) Minstrel: a computer model of creativity and storytelling. PhD dissertation, University of California, Los Angeles
31. Ware SG, Young RM (2014) Glaive: a state-space narrative planner supporting intentionality and conflict. In: Proceedings of the 10th AAAI conference on artificial intelligence and interactive digital entertainment. AAAI, Menlo Park
32. Yannakakis GN, Togelius J (2011) Experience-driven procedural content generation. IEEE Trans Affect Comput 2:147–161
33. Zillman D (1996) The psychology of suspense in dramatic exposition. In: Vorderer P, Wulff HJ, Friedrichsen M (eds) Suspense: conceptualizations, theoretical analyses, and empirical explorations. Lawrence Erlbaum Associates, Mahwah, pp 199–231
34. Zwaan RA, Langston MC, Graesser AC (1995) The construction of situation models in narrative comprehension: an event-indexing model. Psychol Sci 6:292–297