Building an Intelligent, Authorable Serious Game for Autistic Children and Their Carers*

Kaśka Porayska-Pomsta[1], Keith Anderson[2], Sara Bernardini[3], Karen Guldberg[4], Tim Smith[5], Lila Kossivaki[4], Scott Hodgins[6], and Ian Lowe[7]

[1] Institute of Education, London Knowledge Lab, 23-29 Emerald Street, London WC1N 3QS
K.Porayska-Pomsta@ioe.ac.uk
[2] Tandemis Limited
[3] King's College London
[4] University of Birmingham, School of Education
[5] Birkbeck College, Department of Psychological Sciences
[6] Acuity ETS Limited
[7] Topcliffe Primary School, Birmingham

Abstract. This paper introduces the SHARE-IT project, which leverages serious games paradigm to motivate and engage children with autism diagnosis in interactive activities, based on the state-of-the-art autism intervention practices. The aim of SHARE-IT is to formulate, in partnership with schools, parents and industry, the requirements for a robust, intelligent and authorable environment for supporting children in exploring, practicing and acquiring social interaction skills. SHARE-IT focuses on two key challenges: (i) developing robust system architecture and implementation, able to support both continuing development of a serious game for children with autism and its real world use; and (ii) selecting appropriate technologies and techniques to allow for (a) multi-device and operating system deployment, (b) the development of an *intelligent* serious game for supporting social interaction while (c) allowing the flexibility for the environment to be authored by lay persons. SHARE-IT's architecture is presented and several considerations of importance to enabling the engineering of an intelligent and authorable serious game are discussed. Examples of technologies developed to date are given throughout and a discussion of future challenges offered.

1 Introduction

Autism Spectrum Conditions (ASCs) are neuro-developmental conditions affecting an increasing number of individuals globally, with conservative estimates reporting approximately 6 per 1000 autistic children under 8 years [1]. Children with ASCs have marked difficulties in social interaction and communication and

* This project is funded by the EPSRC, grant number: EP/K012428/1. We thank Cathy Ennis and Arjan Egges (Univerity of Utrecht) for their advice about agent animations and for a set of motion-captured animations for our agent.

D. Reidsma, H. Katayose, and A. Nijholt (Eds.): ACE 2013, LNCS 8253, pp. 456–475, 2013.

in initiating and responding to social actions, including imitation, turn-taking and collaborative (joint) actions. Early intervention and consistent support that is also sustained over time and contexts is paramount to improving children's ability to cope with social situations and to enhancing their and their caregivers' quality of life [2]. Modern interventions emphasise the need for consistent support across contexts, and for teachers and parents to share the management of goals for each child through co-creation of learning experiences [3].

Increasingly, teachers and parents look to technology as a complimentary support. Unfortunately, there is a notable paucity of autism related applications that would allow parents or teachers to author and modify the existing software according to the specific needs of individual children, without reliance on software developers. Furthermore, despite the growing number and availability of technology-enhanced interventions (TEIs) for autism, as with previous teaching-learning innovations, design and research have evolved with little direct influence on practice and limited research in real-world classrooms and homes. Many TEIs are research prototypes, which are rarely robust enough for a researcher-free deployment *in-the-wild* and they often lack substantial evaluations to demonstrate their role and efficacy as intervention tools [4].

SHARE-IT's aim is to formulate, in partnership with schools, parents and industry, the requirements for a robust, intelligent and authorable serious game for supporting children with ASCs in exploring, practicing and acquiring social interaction skills. SHARE-IT builds on the ECHOES project, in which a serious game for children with ASCs was designed and evaluated [5]. ECHOES facilitated interactions through a multitouch screen and it employed a semi-autonomous agent, which was underpinned with AI techniques, including automated planning and user modelling for real-time tracking of the child's progress.

The outcomes of ECHOES motivate SHARE-IT: (1) ECHOES provoked children to manifest social behaviours that they did not exhibit in the classroom and observing such behaviours allowed teachers to appreciate the child's hidden potential and to tailor the support accordingly; (2) it demonstrated a potential to improve key skills for social communication, such as turn-taking, sharing of attention with others, and initiating and responding to bids for interaction; and (3) highlighted the real demand for affordable TEIs for Autism [5]. ECHOES also highlighted several key design and engineering challenges, specifically: (4) the need for robust system architecture and implementation, able to support both continuing development of a serious game such as ECHOES and its real world use; (5) the need for considered selection of appropriate technologies and techniques to allow for (a) multi-device and operating system deployment, (b) the development of an *intelligent* serious game needed in the context of supporting social interaction while (c) offering the flexibility for the environment to be authored by lay persons.

SHARE-IT's chief focus lies in addressing challenges (4) and (5). It targets children's *social communication* competences, including the ability to coordinate and share attention, intentions, and emotions with others, and engaging in reciprocal interaction through verbal and non-verbal means. Crucially, SHARE-IT investigates the question of the appropriate infrastructure that is needed to foster independent and sustainable engagement and flexibility of the serious game use by stakeholders, at the component, device and operating system levels.

The remainder of the paper is structured as follows: Section 2 provides a brief introduction to the trends in serious games for social skills training for autism as well as in *intelligent* games. In Section 3, we present SHARE-IT's pedagogical underpinnings and how these are reflected in the SHARE-IT's activities and design decisions, including the need for autonomous agents and real-time user modelling, which we argue are essential for developing serious games that foster social skills in children with ASCs. In Section 4, we present SHARE-IT's architecture and discuss several considerations of importance to enabling the engineering of an intelligent serious game. Section 5 is dedicated to examining the decisions that are of relevance to making a game such as SHARE-IT authorable by teachers and parents. In Section 6, we offer some conclusions and outline the future work.

2 Background

2.1 Serious Games for Autism

Serious games are increasingly employed as *tools* for supporting a wide range of activities, from therapy, e.g. for pain management and rehabilitation, to training of specific skills, such as literacy and numeracy [6]. Educational games based on virtual reality are believed to offer particular benefits for children with ASCs [7], because they can help alleviate children's stress associated with real social interactions and because they allow the child to repeatedly rehearse behaviours in different scenarios, from simple and structured situations to increasingly more complex and unpredictable ones. Studies show that autistic children who were taught by a virtual human experience higher levels of retention than those in traditional classroom settings [8] and that virtual agents may promote generalisation [9] - a holy grail of any autism intervention. There is also some evidence that both role-play and practice of behaviours across different contexts can contribute to increasing the chances of transferring the learned skill from the virtual to the real world [10].

The majority of serious games for autistic children have focused mainly on the following domains of social communication:

Language Skills. Games targeting these skills range from systems that aim to increase the fluency of speech [11] to systems that improve the intelligibility of speech [12]. Since understanding child speech is still technologically challenging, some of the more successful applications involve a "Wizard of Oz" approach, where a human experimenter manipulates the system [9]. Recent improvements in speech recognition allow for technologically more sophisticated support (e.g. TouchStory [13]), and open up the possibility of focusing on literacy skills and expanding children's vocabulary [14,15].

Affective Skills. The focus of games in this category is on training recognition of facial expressions and body gestures, e.g. cMotion [16] and LIFEisGAME [17], both of which aim to help children recognise contextualised facial expressions through manipulation of an interactive virtual character [18]. Commercial games, include FaceSay [19] or Society [20], which provide children with opportunities to practice key skills, e.g. attending to eye gaze, discriminating facial expressions and recognising faces and emotions with the help of interactive avatar assistants, or which teach how to recognise and control their emotions and cope with social challenges such as handling bullying and talking to strangers.

Interaction Skills. Turn-taking, imitation and collaborative play provide the focus for games in this category. The tabletop interactive Collaborative Puzzle Game [21] aims to foster collaboration skills. It revolves around an interaction rule, called "enforced collaboration", where puzzle pieces must be touched and dragged simultaneously by two players in order to be moved. Barakova et al. [22] proposed a multi-agent system composed of autonomous interactive blocks that can express emergent behaviours through change in their colours and light intensity based on how they are manipulated by the users. Teaching interaction skills to autistic children through robotic toys is an area of emerging interest. Examples include a silicone robot, called Keepon [23] and different robotic platforms used in the AuRoRa project [24], from simple mobile robots to more anthropomorphic creatures.

2.2 Intelligent Serious Games

With the growing sophistication of user expectations, a demand emerges also for more convincing game characters and plots, which can be adapted in real-time based on the knowledge that the system can acquire about the user behaviours, also in real-time. Artificial Intelligence techniques and paradigms increasingly provide a research and application focus in this context, with autonomous agents and user modelling emerging as two areas of special interest to games design.

Autonomous Agents. Social interaction involves social partnership. This creates a demand for characters which can act as such partners to the user in a believable way. We argue that credible social partnership requires an investment into the creation of autonomous or at least semi-autonomous agents, i.e. agents that are able to decide independently how to act best to achieve a set of high-level goals that have been delegated to them. Autonomy involves a broad spectrum of behaviours with *no autonomy* and *full autonomy* lying at its extremes. Autonomous agents carry a significant potential for autism intervention because they can contribute to the intensive one-on-one support needed while easing the demand for such support from practitioners and parents. Such agents can complement the traditional interventions allowing human practitioners to focus on the most complex aspects of face-to-face interventions, while supporting repetitive tasks and on-demand access.

Although in the last ten years there has been a growing interest in the potential of artificial agents in the area of autism intervention, both virtual and physically embodied, the efforts have focused primarily on agents with little or no *autonomy*. Typically, in such contexts, virtual agents are either authored *a priori* or controlled by a practitioner through a control panel, while robots are tele-operated. The approaches which form the exceptions in this respect include the Thinking Head project [25], which focuses on developing a talking head that teaches social skills, through realistically portraying facial expressions, and the virtual peers, Baldi and Timo, [14] – 3D computer-animated talking heads for language and speech training. The limited effort in the area of autonomous agents for autism intervention can be in part explained by the fact that, in line with some behaviourist clinical intervention frameworks, most technologies in this context focus on *training* children with respect to specific skills, e.g. recognising a predefined set of facial expressions, rather than on creating believable social interaction experiences, which can be freely explored by the child.

Planning techniques are crucial to enriching virtual characters with autonomy. Path planning, for example, is extensively used in action games to avoid collisions between multiple characters and to allow characters to inhabit and move around architecturally complex environments [26,27]. Task planning is used to enhance the quality of individual moves in puzzle games [28] and to help coordinate actions by multiple characters [29]. Narrative and storytelling planning are of particular relevance to serious games for supporting social interaction, because their aim is to support the generation of evolving contexts and stories within which characters are able to manifest intentionality and emotions [30]. Such emergent narratives are useful in contextualising the behaviours of characters and the training of social skills as reviewed in Section 2.1

User Modelling. In general, user modelling consists of: (1) static profile of the user, containing information such as gender, age, preferences etc., and (2) a diagnostic system responsible for inferring the user's hidden psychological states from behaviours observable online. Specifically to games, recent reviews discuss modelling players' behaviours in terms of their actions vis á vis the games' objectives, their tactics (short-term behaviours) and strategies (long-term action planning) and player profiling in terms of their psychological predispositions based on the detected actions, tactics and strategies [31]. In this context, the player modelling applications tend to focus on adult players in shoot-them-up or strategy games. Relatively little effort is available in relation to modelling children as players, with the notable exception of work by [32], who create and evaluate user models that predict 8-11 year old children's entertainment of games.

The recent recognition of serious games as a genre with substantial potential in complimenting traditional education motivates the need for user modelling in the context. Much work that is relevant to modelling players of serious games comes from the related fields of Intelligent Learning Environments and Affective Computing. Many accounts exist of the nature of mental states that relate to user motivation and learning. Malone [33] highlights factors such as *challenge*, *curiosity* and *fantasy* that impact on users' enjoyment of and learning with digital environments and these factors are also increasingly investigated in relation to users of serious games. Csikszentmihalyi [34] discusses *flow* as crucial to users' ability to immerse themselves in virtual worlds, with the specific dimensions of flow in relation to children game playing, being identified as users' *endurability*, *engagement* and *expectations* [35]. Flow refers to users' sustained focus of attention which psychologists link to children's improved memory, learning and increased ability to self-regulate emotionally [36]. Burleson [37] discusses the state of *stuck* as a factor potentially detrimental to users' enjoyment and learning. User *achievement of the goals* within a specific environment and the user's *personal goals* that emerge from a combination of their mental states and personality traits was examined by Ortony et al. [38] and most notably implemented in a learner model of 6th and 7th grade students (in the US) by Conati [39]. Research in Artificial Intelligence in Education provides further insights into the older children's mental states that are important for learning with technology (e.g. [40]), how these may relate to the specific forms of digital environments and contexts of use (e.g. [41]) and types of help provided to the child (e.g. [42]).

Research in affective computing and use of sensors for detecting users' behaviours in games and education, is thriving. Most existing approaches rely on a combination of cues, such as touch, eye-gaze, facial expression tracking and voice recognition [43]. A recent example of user modelling that relies on cue detection in the context of serious games for social interaction include the work by the TARDIS project, which focuses on developing a technology for training young adults at risk of social exclusion to self-regulate emotionally and to self present to potential employers [44].

3 SHARE-IT's Pedagogical Underpinnings

SHARE-IT focuses on enhancing the social communication competences of children with ASCs, because this is the domain with which they typically have the most difficulty [3]. Social communication involves the ability to coordinate and share attention, intentions, and emotions with others and a capacity for engaging in reciprocal interaction by understanding and using verbal and non-verbal means. SCERTS [3] provides SHARE-IT with a comprehensive approach to social communication assessment and intervention in autism. Based on several established methods in clinical and educational interventions, it identifies the particular skills that are essential for successful social communication: (i) *Social Communication*: spontaneous and functional communication, emotional expression, and secure and trusting relationships with others; (ii) *Emotional Regulation*: the ability to maintain a well-regulated emotional state to cope with stress and to be available for learning and interacting; (iii) *Transactional Support*: caregivers' ability to respond to the child's needs and interests, to adapt the environment, and to provide tools to enhance learning.

3.1 SHARE-IT's Learning Activities

SHARE-IT serious game is organised around discrete scenes, each with a definite pedagogic purpose. For example, a child might be encouraged to turn-take to grow flowers by shaking the cloud (Fig. 1) or stack pots to build a tower (Fig. 2).

SHARE-IT uses the design of the ECHOES learning activities, which have been created over several years through participatory design workshops with teachers, practitioners and children, e.g. [5]. The activities are set in a sensory garden, which is populated by an agent, called Andy, and by interactive objects that react in different ways, sometimes transforming into other objects when the agent or the child act upon them through specific touch gestures or through gaze. For example, tapping on a flower can turn it into a floating bubble or a bouncy ball. The child can pop the bubbles simply by looking at them. The relationship of action-reaction implemented through touch and gaze is intended to support children's understanding of cause and effect, which is often impaired in individuals with ASCs, while the gaze controlled interface also supports children's development of attentional control.

The choice of a magic garden as the setting for the learning activities serves a number of purposes - mainly encouraging the child's imagination through the unusual properties of the objects which can do things that are not possible in the real world, and through exploration, since it is not obvious to the children from the outset how the different objects will behave and react to their gestures. In addition, the use of the magic garden builds on the SCERTS principle that learning activities need to share an obvious unifying theme in order to support shared attention [3].

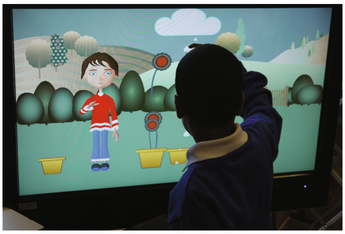

Fig. 1. A child interacting with the ECHOES system through multi-touch interface. In SHARE-IT, we reuse the types of the activities and most of the skin (appearance) of the objects and the environment.

Fig. 2. Andy, the SHARE-IT agent demonstrates pot stacking to the child. The idea is that the child will imitate Andy's action and take turns with Andy in subsequent moves.

SHARE-IT's agent, Andy, acts as a social partner to children, prompting them to take their turn. This requires the agent to be aware of when the child has completed an action and to be capable of taking their own turn. In addition the agent needs to be able to assess the appropriateness of the action for the activity and with respect to the agent's own actions. Finally, Andy needs to deal with the unexpected, e.g. the user not responding, not taking turns, or interacting with the environment in ways unrelated to the activity. Unlike in many games, the difficulty levels in SHARE-IT do not correspond to the different activities, but rather they are organised around the different types of pointing, from relatively ambiguous distal pointing with gaze (i.e. looking at an object), to proximal pointing by touching an object of interest. A combination of several types of pointing, such as distal pointing through gaze and finger point, as well as the agent's verbal request, e.g. "Pass me that pot" is considered the least difficult in SHARE-IT. Prompting and verbal guidance, such as Andy reminding the child of their turn is also used to regulate the difficulty level within activities. Children who respect turn-taking will get fewer prompts, whereas more scaffolding is needed for children who have difficulties in this area.

4 Engineering SHARE-IT's Serious Game

SHARE-IT is essentially a complex system (Fig. 3), combining multimodal user data, specifically – screen touch and eye-tracking data, with facial recognition data being planned as additional input, and requiring for all of these data to be routed to its relevant subsystems. The primary subsystems are the planner and the user model and both of these are consumers of the user data. From a purely engineering point of view, of importance to building robust complex systems for real world use, this leads to questions of synchronisation and prioritisation, specifically – which subsystem should receive the inputs first and then how should the outputs from that subsystem be combined with the results of feeding the data to the other subsystems?

Many systems such as ECHOES are based on a distributed architecture using middleware like ICE [45] to handle inter-process communication. This approach allows for different languages to be used (e.g. ECHOES used Java and Python primarily) and, depending on the languages chosen, it can provide cross-platform

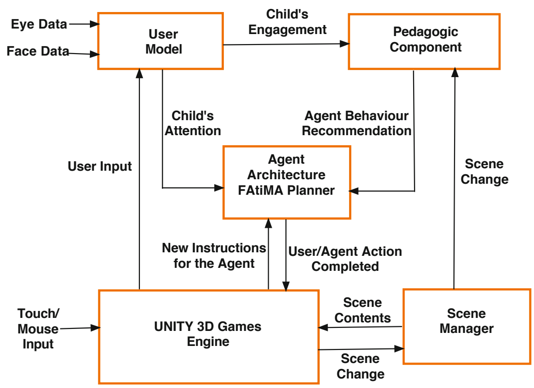

Fig. 3. SHARE-IT's system architecture

support. However there are major drawbacks with this approach, which can hinder long-term extendibility of the system, its computational efficiency and consequently its real world robustness. Specifically, developing a system that relies on a distributed architecture is very resource heavy, because a separate process is needed for each component (and possibly Java Virtual Machine). Furthermore, debugging such an architecture is difficult, because a separate debugger is needed for each component. Based on our experience and evaluation of the ECHOES system, we opine that a distributed architecture in this context increases system complexity without adding any significant gain.

The idea behind SHARE-IT is to deliver a platform for authorable game development and this, coupled with the fact that there is already significant complexity involved in implementing an autonomous agent, leads to our adopting the philosophy to keep the SHARE-IT system architecture and code design as simple as possible at every level, with additional layers of abstraction being avoided wherever possible. To this end, the whole system is designed around a single process, using just one language. All system resources are kept in one location avoiding the problem of contention as all threads in the system are managed from the one process. Additionally, using a single language allows us to maintain a high degree of code consistency making the project easier to maintain and extend.

Another decision relates to the use of an existing games engine to generate the SHARE-IT game. While the system could have been written from scratch, as indeed it was for ECHOES, a lot of unnecessary work and complexity is added to the project for no appreciable gain. There are many commercial and freeware games engines that can handle all the rendering, audio file handling, character animation, 3D effects and physics that are required for projects like SHARE-IT. Any games engine of this type will automatically handle user input from mouse and touch screen. However, eye-tracking and facial expression recognition need to be handled separately. Since the User Model is the primary consumer of these data, which it uses to make judgements about the child's level of engagement with the system and the frequency and consistency of the desired behaviours, these inputs are routed directly to it. The output from the User Model is then routed to the Planner to allow it to make more informed decisions about how the agent should collaborate with the user in the task at hand. In addition to the User Model, the Pedagogic Component also communicates with the Planner providing a top level control over what forms of support should the agent provide to the individual children along with the frequency and the kind of prompting, thereby being responsible also for setting the difficulty levels for each individual child. The Pedagogic Component is essentially a rule-base expressing SCERTS recommendations, which it executes based on information about the child's engagement supplied by the User Model.

4.1 The Game Engine

At the heart of an agent based game is the agent itself. While such a game requires an environment, this generally consists of background artwork and foreground 3D objects. Although the objects may be animate, they are so in very simple ways compared to an agent. An agent itself comprises a physical embodiment: a number of manifest behaviours – generally animated movement and speech, and rules to control these behaviours.

The choice of artwork to be used for the agent and the environment is critical for the success of a game. However life-like the agent's behaviour or enticing the activities, the child will not be drawn in unless the physical environment of the game is comfortable and attractive. Ideally, the appropriate art design should involve a participatory effort between children, practitioners and designers to ensure that the design is attractive to children, while lending itself to supporting the pedagogical aspects of the game. For example, in activities focused on improving children's attentional control, it is important that the number, the type and the interactive nature of objects acting as distractors are in line with the pedagogic requirements of the task. This paper does not address the creation of artwork, since its focus is on game development. In SHARE-IT, we re-use the artwork that has been inspired by the children and teachers themselves, and which proved very successful with most of almost a hundred children who used the ECHOES environment over the years.

However, it is important to consider that in the context of serious games for social interaction, there is an increased need for precision in the execution of certain gestures of the agent and facial expressions, especially if the expectation is for the child to recognise emotions in such expressions, follow pointing by the agent or imitate the behaviours modelled by the agent. Large libraries of resources are available for choosing artwork, for example TurboSquid[1] and RocketBox[2] and both provide rigged characters that can be animated for use as agents. Motion capture offers the best results, especially for fine-grained hand gestures such as one-finger pointing needed by an agent to unambiguously point to objects, e.g. in initiating bids for interaction with the child, as well as Makaton sign gestures, often used to support instructions and communication with children with autism (see Fig. 4).

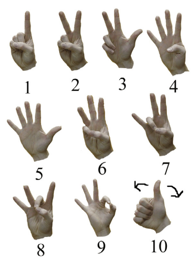

Fig. 4. Makaton numbers gestures illustrate detail that may be needed

[1] `http://www.turbosquid.com/`
[2] `http://www.rocketbox-libraries.com/`

To address the multi-platform requirement of SHARE-IT (Windows, Mac OS, IOS and Android), needed to facilitate access to the SHARE-IT technology by different stakeholders and in different contexts, various games engines were evaluated. Some open source engines, like Ogre[3] are graphics engines only, others like Panda[4] are full games engines, but do not support the range of operating systems required. There are a bewildering number of commercial games engines available (Game Maker Studio, LibGDX, Unity, Blender, Leadwerks, Source, CryEngine). Many do not support all platforms and most are designed for the creation of shoot-them-up style games. Unity 3D was chosen on the basis that it supports the four platforms of interest and that it comes with Mono built-in, which means that JavaScript or C# can be used for the development. The latter is the obvious choice for the complex system such as SHARE-IT, because it offers a fully object-oriented development language rather than a scripting language. Although some would argue that this choice is a matter of preference, we argue that this decision impacts the ease of maintaining the large and complex code-base required by the system like SHARE-IT and which would have been very difficult to achieve with JavaScript.

Another reason that led to the choice of Unity concerned the source of the 3D characters and animations required for the agent. There are a number of systems available (3ds Max, Maya, SketchUp, Realsoft3D, Cheetah3D), but few offer as complete end-to-end support for creating and then animating characters as the AutoDesk products. Furthermore the 3ds Max biped used in 3DS Max is one of the main industry standards and off-the-shelf characters rigged with this biped are widely and cheaply available. There are large libraries of motion capture data available, such as the Carnegie Mellon repository,[5] which can be readily imported onto the 3DS Max biped, and AutoDesk MotionBuilder product provides rich support for the import of motion capture and animation data of different formats. The resulting animations can be exported (using the FBX file format) so that they can be directly imported into Unity3D. All this provides a seamless workflow that further justifies the choice of these technologies.

4.2 SHARE-IT's Intelligence

Agent Development. SHARE-IT employs FAtiMA [46] domain-independent architecture to underpin its agent. FAtiMA combines the type of *reactive* and *cognitive* capabilities that are needed to implement an autonomous, *socio-emotionally* competent agent that is also proactive and reactive. The cognitive layer of FAtiMA is based on *planning* techniques, specifically a variant of PDDL2.1 [47], and it is characterised by: (i) a set of internal goals; (ii) a set of action strategies to achieve these goals; and (iii) an affective system, while the emotional model is derived from the OCC theory of emotions [38] and the appraisal theory [48]. A FAtiMA agent builds an affective system from emotional reaction rules, action tendencies, emotional thresholds and emotion decay rates.

[3] http://www.ogre3d.org/
[4] http://www.panda3d.org
[5] http://mocap.cs.cmu.edu/

Appraisal and *coping* are mechanisms controlling a FAtiMA agent, which is able to "experience" one or more of the 22 emotions of the OCC model, based on its appraisal of the current state of the world and its subjective tendencies to experience certain emotions instead of others. The agent deals with these emotions by applying problem-focused or emotion-focused coping strategies. Both the appraisal and the coping work at the *reactive* level, which affects the short-term horizon of the agent's behaviour, and at the *deliberative* level, which relates to the agent's goal-oriented behaviour.

Similar to ECHOES, in SHARE-IT, each learning activity has an associated FAtiMA agent model. All these models share the same specification of the agent's affective system, to enable the agent to maintain the same 'personality' between sessions, in order to establish a trusting relationship with the child. Andy was designed to have a positive, motivating and supportive character, to be happy and to not get frustrated easily. Such behaviours were obtained through manipulating Andy's emotional reaction rules, thresholds and decay rates of the OCC emotions available. For example, the child looking at an object pointed to by the agent has a very high desirability to it, whereas the child looking away has low desirability. Currently the agent's happiness threshold is set to low, its frustration threshold is high, its happiness decay is slow, while its frustration decay is fast, leading to the agent's tendency to be happy and to not become easily frustrated regardless of the child's actions. Although ECHOES' evaluation, especially the analysis of children's engagement showed that children did interact with Andy as if with an intentional being [49], it also highlighted a possible need to make our agent more interesting, by furnishing it with more moods and emotions, both positive and negative. Our current discussions with teachers and parents are consistent with this analysis, suggesting that a more modulated behaviour is required of the agent within activities to make him more credible and enticing to children, as well as to introduce different levels of difficulty for children, whereby the child is exposed to unexpected and varied behaviours on the part of the social partner. For example, in an activity where the child can shake the cloud to make rain and to grow flowers, children often flout the implicit expectation that cloud ought not to be shaken on the agent. Yet, naturally, children find trying to make the agent wet very funny. Currently, the agent's reaction to this is to smile and continue in good humour. However, following teachers' assessments, by not having the agent progressively become sad or angry, is a missed opportunity for demonstrating to the child the cause-effect relationship between their actions and emotional reactions of others.

Following teachers' suggestions, in SHARE-IT we are developing the agent as an emotional being able to express complex emotions. To this end we have developed authoring tools (discussed in Sec.5; see Fig. 5) for displaying the agents emotions and are in the process of linking those to the more complex specification of its action tendencies in the FAtiMA model, which then control the agent's facial expressions and gestures. With the need for a more sophisticated real-time emotional display by the agent, we are presently working with teachers and parents on articulating: (i) the high-level *pedagogic goals* on which the individual

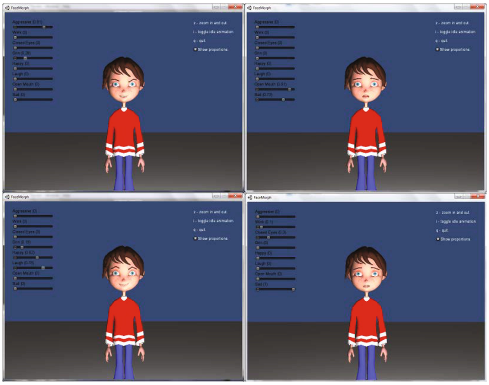

Fig. 5. SHARE-IT authoring tool for morphing the agent's emotional expressions

activities ought to focus and (ii) the specific *narrative content* of the activities itself. These will be incorporated in the Pedagogic Component of the system and will be used to inform the goals for the agent to actively pursue along with appropriate action strategies for each learning activity.

User Modelling. SHARE-IT also reuses ECHOES' user model, which estimates the child's cognitive and affective states in real-time and feeds back this information to the Pedagogic Component and the Planner whenever a change in a given state is detected. In SHARE-IT, the model receives real-time information from the touch and eye-gaze systems and, based on this data, it interprets the current level of child's engagement as well as the level, consistency and desirability of the child's attention vis á vis the goals of a given activity or specific agent actions. The cognitive assessment is facilitated by a *rule-based engine* which estimates the extent to which the child has achieved the goals associated with the session, e.g. based on the number of the child's responses to the agents bids for interaction in a pre-specified time frame. The rules are based on SCERTS guidelines and precise timing constraints for establishing the child's mastery of joint attention and symbolic use skills. For example, the behaviour "shifts gaze between people and objects" is satisfied if the child shifts gaze spontaneously between a person and an object at least three times and the entire sequence occurs within two seconds. Detailed real-time eye-gaze data is essential to enable the interpretation outlined.

In ECHOES, the assessment of the child's engagement was based on a combination of supervised and unsupervised learning techniques [49]. Engagement is defined in terms of: *very engaged*, *engaged* and *not engaged*. The children are considered "very engaged" when they are fully absorbed by the learning activity; "engaged" when they are interested in the current learning activity, but are not immersed in it; "not engaged" when they do not interact with the system at all, either *via* gaze or touch.

In SHARE-IT our efforts to date focused on practicalities of improving the reliability of gaze detection to enhance the estimation of children's engagement. Specifically, we have incorporated a portable Tobii eye-tracker (X2-30 compact) with the rest of the system. Tobii provides us with high precision of gaze detection in different lighting conditions and can cope with different face occluders, such as glasses. Another advantage of this technology is that it comes with an SDK that supports .Net version 2 allowing for a direct integration into a Unity build. SHARE-IT's User Model provides a callback function that gets called by the eye-tracker each time there is new data.

One challenge involved in using eye-trackers in this context relates to calibration. This typically consists of drawing calibration objects on the screen that the user has to stare at in the order they are displayed. Keeping track of our ambition to create a game for children, in SHARE-IT we turned the calibration task into a game-like activity, where the agent instructs and prompts the child where to look. However, there is still the possibility that the child may be unable or unwilling to follow the process, in which case the eye-tracker informs the system that the calibration was unsuccessful. In this case, the process can be retried, or this form of input to the system can be disabled. Since it is envisaged that the system will run on a variety of systems, where eye-trackers may not always be available, the use of this data needs to be optional to the system. The system was thus designed in such a way that the eye-tracker is used to refine the estimates of engagement reached from processing input from mouse/touch and facial expression processing.

The SHORE object recognition engine will be used in SHARE-IT for simple image recognition based on the output from a webcam, such as the built-in camera that features in laptops and tablets, to extract features such as facial locations and a range of emotional expressions. The engine is reasonably good at detecting happy (smiling), sad (downturned mouth) and surprised (open mouthed) expressions, without the need for over-exaggeration [44]. This will provide another source of useful input to the User Model, for example an unhappy looking user may require more intervention and support from the agent. SHORE's SDK supports Mac/IOS and Windows, but not Android, so again this input will be made optional to the system. The data produced by SHORE is of a less definite nature than the other inputs - mouse/touch and eye-tracking data can tell us exactly which objects, or parts of the scene are currently engaging the user. Facial expressions need to be collated over time to provide meaningful data. We are currently developing a suite of activities that will rely on SHORE, involving mimicking the agents' expressions (as can be produced in the tool shown in Fig. 5) and in turn having the agent to mimic the child's facial expressions. Apart from the logs of this data contributing to the training of emotion interpretation capabilities of our User Model, they are also intended to support important skills, such as recognition of emotions from other people's faces and imitation skills.

5 Authoring Tools

The authoring environment for SHARE-IT is intended for use by teachers and parents, so that they can tailor the game according to the needs of the individual children in their care. The authoring environment is designed to allow existing activities to be chosen and new ones to be created from a palette of components. A great deal of effort in constructing the Unity game for the project was spent creating a range of objects and agent animations and speech to provide the building blocks for this tool. Ease of use is a primary consideration given the non-specialist nature of the end users and is a continuing focus of participatory design workshops with teachers and parents.

The authoring application consists of three components: (1) **A scene selector tool** (Fig. 6) that allows the scenes and their order to be selected. It also controls display of a scene selection menu in the game which permits the user free choice of scenes. Furthermore, basic object properties, such as object transformations can be set from this too; (2) **A scene creator/editor window** (Fig. 7), which is used to setup or alter the physical components of a scene - how many of each object and where they are placed, their colour and size, and the scene background; (3) **A rule creation dialogue** (Fig. 8), which is used after the scene creator to set up the planner rules which will govern the agent's behaviour in the scene.

The first two of these components write the XML initialisation file for Unity to read on startup. This allows Unity to create all the required scenes at runtime.

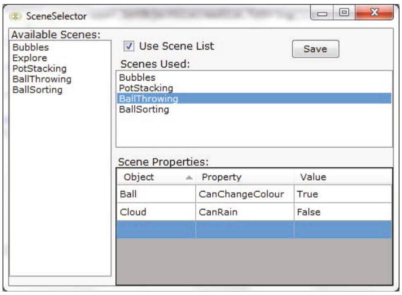

Fig. 6. SHARE-IT authoring tool for selecting and ordering learning activities

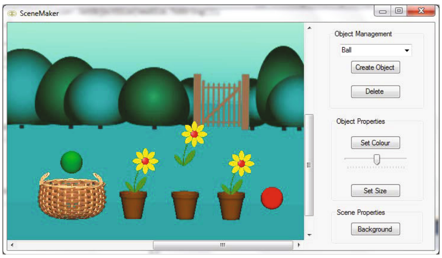

Fig. 7. SHARE-IT authoring tool for choosing the objects, their properties and number

The rule creation component is somewhat more involved, as it has to map the elements selected by the scene author into rules that the FAtiMA planner can consume. As shown in Fig. 8 this might be an activity to throw bouncy balls through the cloud to turn them different colours and in which the agent and user take turns with each ball. Each action type has an XML template that forms the basis of the rule in Fig. 9, which then gets filled in with the details of the objects

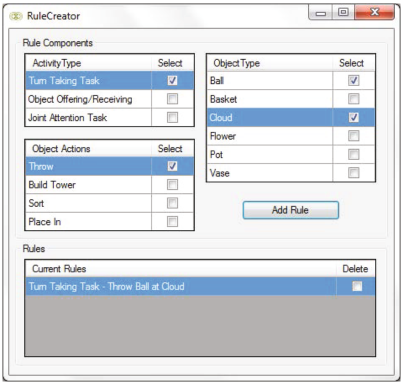

Fig. 8. SHARE-IT authoring tool for creating planner rules.

```xml
<Action name="SelfThrowBallThroughCloud([sourceObject],[targetObject])" probability="1">
  <PreConditions>
    <Property name="BallThrowing(isChosenActivity)" operator="="  value="True" />
    <Property name="indicatedTurn()" operator="=" value="True" />
    <Property name="[targetObject](type)" operator="=" value="Cloud" />
    <Property name="[sourceObject](type)" operator="=" value="Ball" />
    <Property name="[sourceObject](beenThroughCloud)" operator="!=" value="True" />
    <Property name="hold([SELF],[ball])" operator="=" value="True" />
  </PreConditions>
  <Effects>
    <Effect probability="1.0">
      <Property name="giveSelfFeedback()" operator="=" value="True" />
    </Effect>
    <Effect probability="1.0">
      <Property name="[sourceObject](beenThroughCloud)" operator="=" value="True" />
    </Effect>
  </Effects>
</Action>
```

Fig. 9. XML rule relating to the rule created in the editor shown in Fig. 8

selected. In the example in Figs. 8 and 9, the "indicatedTurn()" property means that this is a turn taking activity and "BallThrowing(isChosenActivity)" means that this rule only applies to the Throwing activity.

Mapping of user friendly representations to the rules governing the planner required a simple set of rule categories to be chosen such that an automatic translation mechanism could be set up. This set needed to be rich enough to allow for the pedagogic aims of the system to be expressed. The rule base for the SHARE-IT planner consists of a number of fixed rules which govern the low-level behaviours of the agent that are common to all activities: (1) Rules that govern the low-level behaviours of the agent, e.g. types of movement and interaction strategies, e.g. whether and when to prompt, correct or encourage the child. These are part of the psychological makeup of the agent and are currently hard-coded. (2) General scene-related activities, e.g. walking on/off scene, starting an activity, detecting and acknowledging the end of an activity, manipulating objects (throwing, stacking, etc.). These are common to all activities and are also hard-coded. (3) Higher-level behavioural traits of the agent, such as types of greeting, speech and movement used for the prompting and encouragement and one-off behaviours, such as giggling or falling around, designed to relax or entertain the child, and not linked to any particular activity.

SHARE-IT authoring tools allow for a specification of higher level rules that determine how the agent behaves during a particular activity. For example, an activity that involves turn taking in stacking objects to build a tower would involve specifying the object type, the action permitted (stacking) and the fact that turn-taking is required. In order to keep the rule base manageable and to prevent the authoring tool itself from becoming too complex, the activities were constrained to the following categories: (i) Turn taking activities, (ii) Object offering/receiving, e.g. allowing the user to hand an object to the agent, for instance, (iii) Joint activities, e.g. carrying an object together. For each of these categories XML templates have been created that can be modified to construct individual activities as specified by the user in the Rule Creation tool (Fig. 8). In this way any combination of objects and actions on objects can be specified to create an activity.

6 Conclusions and Future Work

In this paper, we introduced the SHARE-IT project, which aims to design and implement an infrastructure for parents and teachers to become co-creators of their children's technology-enhanced learning experiences. We outlined a number of practical considerations of importance to building intelligent and autonomous serious games for children with autism and for their carers. The work reported is very much in progress, albeit it is based on substantial previous work related to the ECHOES project, which it extends. The considerations are therefore grounded in our extensive experience of developing an intelligent serious game for social interaction. Much work still remains to be done, particularly in relation to enriching our virtual character with more interesting and varied behaviours, extending the range of characters available to include females and those representing different ethnicities. A major limitation in the agent's intelligence is that the agent cannot currently deal with inappropriate or unexpected behaviour from the user, such as not responding to the agent, not honouring turn taking, being occupied with objects in the scene not related to the activity, complete disengagement, etc. We are currently adding further rules to the planner's actions to cover the most obvious cases, although in the longer term it would be desirable to make these behaviours part of the agent's personality, so that they would vary with the affective system parameters. We are also focusing on improving the accuracy of the SHARE-IT modelling tools. Incorporating a robust eye-gaze tracking system is of special importance in this context because the ability to follow gaze by social partners provides the foundations of successful social communication. While the Tobii technologies currently employed are still unaffordable by most, they are set to enter the mainstream market as part of ordinary PCs in the next 12 months. We are currently working with parents and teachers on refining the design and the variation of the SHARE-IT learning activities to make them more flexible to extend and we are testing the authoring tools to assess their usability and robustness. In the next 3 months studies with children, parents and teachers are planned, where the SHARE-IT technologies will be piloted simultaneously in school and home contexts and the authoring tools will be tested in-the-wild, with a special attention being paid to the relevance and feasibility of the activities across the two contexts. While children are the main beneficiaries of the game developed, providing parents and teachers with authoring powers is also expected to enhance parents-school communication and continuity, consistency and coherence of the experience provided to children. The results of the planned studies will be reported in due course.

References

1. Medical Research Council: Autism - Research Review (2001) (accessed October 20, 2012)
2. Knapp, M., Romeo, R., Beecham, J.: The economic consequences of autism in the uk. Foundation for People with Learning Disabilities Report (2007)
3. Prizant, B., Wetherby, A., Rubin, E., Laurent, A., Rydell, P.: The SCERTS® Model: A Comprehensive Educational Approach for Children with Autism Spectrum Disorders. Brookes (2006)
4. Wass, S., Porayska-Pomsta, K.: The uses of cognitive training technologies in the treatment of autism spectrum disorders. Autism: International Journal of Research and Practice (in press)
5. Bernardini, S., Porayska-Pomsta, K., Smith, T.J., Avramides, K.: Building autonomous social partners for autistic children. In: Nakano, Y., Neff, M., Paiva, A., Walker, M. (eds.) IVA 2012. LNCS (LNAI), vol. 7502, pp. 46–52. Springer, Heidelberg (2012)
6. de Freitas, S.: Learning in immersive worlds. Technical report, Joint Information Systems Committee, Bristol (2006) (accessed October 20, 2012)
7. Parsons, S., Cobb, S.: State-of-the-art of virtual reality technologies for children on the autism spectrum. European Journal of Special Needs Education 26(3), 355–366 (2011)
8. Grynszpan, O., Martin, J.C., Nadel, J.: Multimedia interfaces for users with high functioning autism: An empirical investigation. International Journal on Human-Computer Studies 66(8), 628–639 (2008)
9. Tartaro, A., Cassell, J.: Playing with virtual peers: bootstrapping contingent discourse in children with autism. In: Proceedings of the ICLS 2008, pp. 382–389 (2008)
10. Parsons, S., Mitchell, P.: The potential of virtual reality in social skills training for people with autistic spectrum disorders. Journal of Intellectual Disability Research 46(5), 430–443 (2002)
11. Anwar, A., Rahman, M., Ferdous, S., Anik, S., Ahmed, S.: A computer game based approach for increasing fluency in the speech of the autistic children. In: Proceedings of the11th IEEE International Conference on Advanced Learning Technologies (ICALT), pp. 17–18 (2011)
12. Rahman, M.M., Ferdous, S., Ahmed, S.I., Anwar, A.: Speech development of autistic children by interactive computer games. Interactive Technology and Smart Education 8(4), 208–223 (2011)
13. Davis, M., Otero, N., Dautenhahn, K., Nehaniv, C., Powell, S.: Creating a software to promote understanding about narrative in children with autism: Reflecting on the design of feedback and opportunities to reason. In: IEEE 6th International Conference on Development and Learning, Proceedings of the Development and Learning, ICDL 2007, pp. 64–69 (July 2007)
14. Bosseler, A., Massaro, D.: Development and evaluation of a computer-animated tutor for vocabulary and language learning in children with autism. Journal of Autism and Developmental Disorders 33(6), 653–672 (2003)
15. Massaro, D.W.: Embodied agents in language learning for children with language challenges. In: Miesenberger, K., Klaus, J., Zagler, W.L., Karshmer, A.I. (eds.) ICCHP 2006. LNCS, vol. 4061, pp. 809–816. Springer, Heidelberg (2006)
16. Finkelstein, S.L., Nickel, A., Harrison, L., Suma, E.A., Barnes, T.: cMotion: A new game design to teach emotion recognition and programming logic to children using virtual humans. In: Proceedings of the 2009 IEEE Virtual Reality Conference, pp. 249–250 (2009)

17. Abirached, B., Zhang, Y., Aggarwal, J., Tamersoy, B., Fernandes, T., Miranda, J., Orvalho, V.: Improving communication skills of children with asds through interaction with virtual characters. In: 2011 IEEE 1st International Conference on Serious Games and Applications for Health (SeGAH), pp. 1–4 (2011)
18. Schuller, B.: ASC-Inclusion Project (accessed October 20, 2012)
19. Hopkins, I.M., Gower, M.W., Perez, T.A., Smith, D.S., Amthor, F.R., et al.: Avatar assistant: Improving social skills in students with an asd through a computer-based intervention. Journal of Autism and Developmental Disorders 41(11), 1543–1555 (2011)
20. Beaumont, R., Sofronoff, K.: A multi-component social skills intervention for children with asperger syndrome: The junior detective training program. Journal of Child Psychology and Psychiatry 49, 743–753 (2008)
21. Battocchi, A., Pianesi, F., Tomasini, D., Zancanaro, M., Esposito, G., Venuti, P., Ben Sasson, A., Gal, E., Weiss, P.L.: Collaborative puzzle game: a tabletop interactive game for fostering collaboration in children with autism spectrum disorders (asd). In: Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces, ITS 2009, pp. 197–204. ACM, New York (2009)
22. Barakova, E., van Wanrooij, G., van Limpt, R., Menting, M.: Using an emergent system concept in designing interactive games for autistic children. In: Proceedings of the 6th International Conference on Interaction Design and Children, IDC 2007, pp. 73–76. ACM, New York (2007)
23. Kozima, H., Michalowski, M., Nakagawa, C.: Keepon: A playful robot for research, therapy, and entertainment. International Journal of Social Robotics 1(1), 3–18 (2009)
24. Dautenhahn, K., Werry, I.: Towards interactive robots in autism therapy: Background, motivation and challenges. Pragmatics and Cognition 12(1), 1–35 (2004)
25. Milne, M., Luerssen, M., Lewis, T., Leibbrandt, R., Powers, D.: Development of a virtual agent based social tutor for children with autism spectrum disorders. In: Proceedings of the International Joint Conference on Neural Networks, pp. 1–9 (2010)
26. Snape, J., Guy, S., Lin, M., Manocha, D.: Local and global planning for collision-free navigation in video games. In: Proceedings of the 3rd International Planning in Games Workshop, ICAPS 2013 (2013)
27. Jaklin, N., van Toll, W., Roland, G.: Way to go - a framework for multi-level planning in games. In: Proceedings of the 3rd International Planning in Games Workshop, ICAPS 2013 (2013)
28. Do, M., Tran, M.: Blocksworld: An ipad puzzle game. In: Proceedings of the 3rd International Planning in Games Workshop, ICAPS 2013 (2013)
29. Menif, A., Guettier, C., Cazenave, C.: Planning and execution control architecture for infantry serious gaming. In: Proceedings of the 3rd International Planning in Games Workshop, ICAPS 2013 (2013)
30. Riedl, M., Young, M.: Narrative planning: Balancing plot and character. Journal of AI Research 29, 217–268 (2010)
31. Bakkes, S., Tan, C.T., Pisan, Y.: Personalised gaming. Creative Technologies 3 (2012)
32. Yannakakis, G., Hallam, J.: Modeling and augmenting game entertainment through challenge and curiosity. International Journal on Artificial Intelligence Tools 16(6), 981–999 (2007)
33. Malone, T.W.: Guidelines for designing educational computer games. Childhood Education 59, 241–247 (1983)

34. Csikszentmihalyi, M.: The Psychology of Optimal Experience. Harper Perennial, New York (1990)
35. Read, J., McFarlane, S., Cassey, C.: Endurability, engagement and expectations: Measuring children's fun. In: Proceedings of the International Conference for Interaction Design and Children (2002)
36. Safyan, L., Lagattuta, K.H.: Grownups are not afraid of scary stuff, but kids are: young children's and adults' reasoning about children's, infants', and adults' fears. Child Development 79(4), 821–835 (2008)
37. Burleson, W.: Affective Learning Companions: strategies for empathetic agents with real-time multimodal affective sensing to foster meta-cognitive and meta-affective approaches to learning, motivation and perseverance. PhD thesis, Massachusetts Institute of Technology (2006)
38. Ortony, A., Clore, G.L., Collins, A.: The Cognitive Structure of Emotions. Cambridge University Press (1988)
39. Conati, C.: Probabilistic assessment of user's emotions in educational games. Journal of Applied Artificial Intelligence, Special Issue on Merging Cognition and Affect in HCI 16 (7-8), 555–575 (2002)
40. Arroyo, I., Beack, J., Beal, C.R., Woolf, B.P.: Learning with the zone of proximal development with the animalwatch intelligent tutoring system. In: Proceedings of the American Educational Research Association annual meeting, Chicago, IL (2003)
41. Kerawalla, L., O'Connor, J., Underwood, J., du Boulay, B., Holmberg, J., Luckin, R., Smith, H., Tunley, H.: The homework system: using tablet pcs as tools to support continuity of numeracy learning between home and primary school. Educational Media International 44(4), 289–303 (2007)
42. Chen, Z.H., Deng, Y.C., Chou, C.Y., Chan, T.W.: Active open learner models as animal companions: motivating children to learn through interaction with my-pet and our-pet. International Journal of Artificial Intelligence in Education 17, 145–167 (2007)
43. Pantic, M., Rothkrantz, L.J.M.: Towards an affect-sensitive multimodal human-computer interaction. Special Issue on Multimodal Human-Computer Interaction (HCI) 91, 1370–1390 (2003)
44. Damian, I., Baur, T., Gebhard, P., Porayska-Pomsta, K., André, E.: A software framework for social cue-based interaction with a virtual recruiter. In: Proceedings of the 13th International Conference on Intelligent Virtual Agents, Santa Cruz (2013)
45. Henning, M.: A new approach to object-oriented middleware. IEEE Internet Computing 8(1), 66–75 (2004)
46. 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)
47. Fox, M., Long, D.: PDDL 2.1: An extension to PDDL for expressing temporal planning domains. Journal of Artificial Intelligence Research 20, 61–124 (2003)
48. Smith, C.A., Lazarus, R.S.: Emotion and adaptation. In: Vlahavas, I., Vrakas, D., Pervin, L.A. (eds.) Handbook of Personality: Theory and Research, pp. 609–637. Guilford, New York (1990)
49. Bernardini, S., Porayska-Pomsta, K., Sampath, H.: Designing an intelligent virtual agent for social communication in autism. In: Proceedings of the Ninth Annual AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (2013)