



ECHOES: An intelligent serious game for fostering social communication in children with autism



Sara Bernardini ^{a,*}, Kaśka Porayska-Pomsta ^b, Tim J. Smith ^c

^a Department of Informatics, King's College London, London WC2R 2LS, UK

^b London Knowledge Lab, Institute of Education, London WC1N 3QS, UK

^c Department of Psychological Sciences, Birkbeck College, London WC1E 7HX, UK

ARTICLE INFO

Article history:

Received 31 October 2012

Received in revised form 18 September 2013

Accepted 22 October 2013

Available online 30 October 2013

Keywords:

Virtual social partner

Pedagogical agent

Autonomous intelligent agent

Artificial intelligence planning

Autism

Social communication

ABSTRACT

This paper presents ECHOES, a serious game built to help young children with autism spectrum conditions practice social communication skills. We focus on the design and implementation of the interactive learning activities, which take place in a two-dimensional sensory garden, and the autonomous virtual agent, which acts as a credible social partner to children with autism. Both the activities and the agent are based on principles of best autism practice and input from users. Specification guidelines are given for building an autonomous socially competent agent that supports learning in this context. We present experimental results pertaining to the effectiveness of the agent based on an extensive evaluation of the ECHOES platform, which show encouraging tendencies for a number of children.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

This paper presents the design and implementation of ECHOES, a serious game built to help young children with Autism Spectrum Conditions (ASCs) practise and acquire social communication skills. In ECHOES, children interact with an intelligent virtual character in the context of social situations through a 42 inch multitouch LCD display with eye-gaze tracking. The three-dimensional agent, which acts credibly both as a peer and as a tutor, inhabits a two-dimensional sensory garden where interactive objects can change their shape and function when the agent or the child touches them. The interaction between the child and the agent is structured around a number of different learning activities intended for real-world use in schools and at home as part of children's everyday routine.

ECHOES has been developed in the context of an interdisciplinary project [40] that builds on recent progress in a number of traditionally independent research areas, including psychology, artificial intelligence, human–computer interaction, technology-enhanced learning and autism intervention [62,5]. In this paper, we focus on the design of the interactive learning activities and the autonomous virtual agent based on participatory design workshops with practitioners and children as well as the SCERTS framework [64] – a well-established educational intervention approach aimed to support social communication (SC) and emotional regulation (ER) of children with ASCs through appropriately designed transactional support (TS).

The remainder of the paper is organised as follows. Sections 2 and 3 provide background information concerning the autism spectrum conditions, a literature review on serious games for autism and the rationale for using virtual agents to foster social behaviours in children with ASCs. Section 4 presents the SCERTS model of social communication, which provides the

* Corresponding author. Tel.: +44 7590835242.

E-mail address: sara.bernardini@kcl.ac.uk (S. Bernardini).

theoretical foundation for the design of the agent and the learning activities. In Sections 5 and 6, we present the details of the design and the implementation of the ECHOES learning activities and the agent. We conclude the paper with Sections 7 and 8, where we describe the evaluation of the agent within the ECHOES virtual environment and our conclusions, respectively.

2. Autism spectrum conditions

Autism is a spectrum of neuro-developmental conditions that affects the way in which a person communicates with and relates to other people as well as how they make sense of the world around them [42]. Conservative estimates of autism prevalence report approximately 60 children with autism per 10,000 children under 8 [41]. Autism is a *spectrum* of conditions because, despite certain characteristics being commonly identified with the ASCs, the exact degree to which such difficulties affect different individuals can vary significantly.

The three main areas of difficulty, known as the “triad of impairments” [3], include:

- (i) *Communication*: problems with both verbal and non-verbal language (for example, literal understanding of language, problems with turn-taking in conversations, difficulties in interpreting facial expressions as well as gestures, and echolalia).
- (ii) *Social interaction*: problems with recognising and understanding other people’s emotions as well as expressing their own emotions (for example, difficulties in understanding and following unwritten social rules, tendency to seem insensitive, preference to spend time alone, difficulties in seeking comfort from other people).
- (iii) *Patterns of restricted or repetitive behaviours*: problems with adapting to novel environments (for example, presence of unusually strong narrow interests, difficulties in coping with unexpected change, restricted social imagination, difficulties in engaging in imaginative play and activities).

Within the autism spectrum, individuals are commonly referred to as *high-functioning* or *low-functioning* depending on the degree of language delay and verbal intelligence quotient (IQ) [3]. ECHOES may be of benefit to individuals across the spectrum, but has been specifically designed for high-functioning individuals (i.e. those with higher verbal IQ).

The causes of autism are currently unclear, but the condition is believed to have both a genetic and an environmental component. Currently, there is no cure for autism. However, there is a wide range of interventions, i.e. methods for facilitating learning and development, that can alleviate autism symptom severity. Amongst the emerging approaches, serious games have attracted increasing attention in the autism community for their educational potential [25,58]. As detailed in the next section, a large variety of games have been proposed to support children in acquiring skills in all the three main areas of difficulty that characterise autism. Among the triad of impairments, ECHOES focuses on enhancing the *social communication* competence of children with ASCs because this is the domain with which they typically have the most difficulty [63] and because recent studies indicate that individuals with ASCs as well as their caregivers consider support for learning social communication skills as the most desirable feature of technology-enhanced intervention [65]. Social communication involves the ability to coordinate and share attention, intentions, and emotions with others as well as the capacity for engaging in reciprocal interaction by understanding and using verbal and non-verbal means.

3. Serious games for autism

In the past decade, a new paradigm has emerged in the development of digital games. Besides their use as metaphors and simulations of the real-world, digital games have increasingly been employed as *tools* to support a wide range of activities, from therapy (e.g. games used as pain relievers and games for rehabilitation) to the training of specific skills (e.g. literacy and numeracy) [19]. This approach targets particularly patients with medical or mental health conditions, for example burns victims as well as people diagnosed with eating disorders and post-traumatic stress disorder [27,51,78].

One of the most promising applications of serious games to therapy and training is represented by tools for supporting children with autism in acquiring social communication skills. Several studies show that the majority of people with autism exhibit a natural affinity with technology and a positive attitude towards computer-based training [65]. This is primarily attributable to the fact that software programs offer a predictable and structured environment that can accommodate their need for organisational support and their preference for routine and repetitive behaviours [50].

Among all available technologies, educational games based on virtual reality are believed to offer particular benefits for children with ASCs for several reasons [15,57,58]. Individuals with autism find real social interactions stressful and intimidating due to their unpredictable and judgemental nature. Hence, traditional educational settings that involve being part of a classroom and interacting with other students and teachers are often challenging for them. The anxiety linked with social interaction can be mitigated by the use of artificial tutors and peers, which can be programmed to act tirelessly, consistently and positively towards the child regardless of the child’s behaviours.

Artificial tutors can support individualised learning that is of particular importance for children with ASCs [18,68], and which in traditional contexts is resource-heavy, since it involves intensive behavioural interventions requiring between 25 and 40 h of practitioner time per week per individual (e.g. the “Early Start Denver Model” [18] and the Lovaas method [35]). An appropriately designed virtual tutor can meet individual children’s needs and allow them to proceed at their own pace. Studies show that children with autism who were taught by a virtual human have higher levels of retention than

those in traditional classroom settings [26] and that virtual agents may promote generalisation [14,72] – the final goal of any autism intervention.

Finally, virtual reality allows children to rehearse behaviours in role-play situations and to exercise the same skill in different scenarios, from simple and structured situations to increasingly more complex and unpredictable contexts. Both visual and auditory stimuli can be carefully controlled in virtual reality and can be gradually increased as the child becomes more familiar and comfortable with the environment. For example, it is well known that children with ASCs find other people's faces particularly difficult to interpret due to the rapid and nuanced information they convey. The complexity of the facial displays of a virtual agent can be controlled by slowly increasing the number and the intensity of the facial expressions presented to the child. There is some evidence that both role-play and practice of behaviours across different contexts can contribute to an increase in the chances of transferring learned skills from the virtual to the real world [58].

The majority of serious games for children with autism have focused on one of the following three main domains of social communication:

- Language skills;
- Affective skills; and
- Interaction skills.

In the rest of this section, we briefly review the main developments in each of these areas. We conclude the section by focusing on the agent technology for children with autism. Additional references concerning serious games for autism can be found in [5,52,57].

3.1. Serious games for improving language skills

Given that autism is often associated with difficulties in developing and using language, many systems for autism focus on training language-related skills. Some of them tackle *speech* training and range from systems designed to increase the fluency of speech [4] to systems that improve the intelligibility of speech [29,66]. Since both synthesising and understanding child speech still presents serious difficulties from a technological point of view, traditionally interactive systems for children with autism that allow spoken input or output have employed a “Wizard of Oz” methodology, where a human experimenter provides input or output to the system [72]. However, progress has recently been made in this area and a few applications are now capable of automatic voice detection through the use of probabilistic methods [23]. A number of applications revolve around the concept of *narrative*, which is considered a fundamental component of the creation and the communication of meaning in social interaction. Tartaro and Cassell [71,72] proposed the use of conversational agents as authorable virtual peers for collaborative narrative creation, while Davis et al. [17] developed the system TouchStory, a simple fill-the-gap interactive game that seeks to improve children's understanding of narratives by decomposing them into their primitive components and representing these components using pictures. Conversational agents in the form of animated cartoon characters have also been used for teaching *literacy skills* and expanding the vocabulary of children with autism and language deficiencies [14,36]. Interestingly, in the area of literacy, positive results have been achieved not only through the employment of embodied virtual tutors, but also through the use of simpler user interfaces based on pictures and sounds [23,49,76].

3.2. Serious games for improving affective skills

There has been a considerable amount of work in supporting the development of affective skills in children with ASCs through serious games, especially with regard to recognition of facial expressions and body gestures. Individuals with autism have great difficulties in recognising, understanding and responding appropriately to other people's facial expressions, tone-of-voice and body gestures as well as understanding and expressing their own emotions [8]. It is believed that “visual supports may be helpful in expanding and enhancing a child's expressive communication system” [63] because children with autism are visual thinkers, i.e. they are often more effective at processing visual information than other types of information. cMotion [21], for example, is designed to teach children with ASCs how to recognise contextualised facial expressions through manipulation of an interactive virtual character using a visual drag-and-drop programming interface. LIFEisGAME [1] aims to teach children with autism to recognise facial emotions using real-time automatic facial expression analysis and virtual character synthesis. The game offers different modalities to the users, including assessing the expression performed by a character, manipulating the facial expression of a three-dimensional avatar, and physically performing a facial expression coherent with a story told to them. ASC-Inclusion [37] is an ongoing project that aims to assist the children in perceiving, understanding and expressing their own emotions through a combination of hardware (microphones, cameras) and software tools (games, text, animations, video and audio clips). Interactive games are used in ASC-Inclusion to train children to recognise their own, and others', facial expressions, tone-of-voice and body gestures. FaceSay [28], which is a commercial product, provides children with low-functioning autism (LFA) and high functioning autism (HFA) with opportunities to practice attending to eye gaze, discriminating facial expressions and recognising faces and emotions in a structured environment where they interact with realistic avatar assistants. The related experimental results are encouraging: the children with LFA demonstrated improvements in emotion recognition and social interactions, while the children with HFA demonstrated

improvements in all three areas: facial recognition, emotion recognition, and social interactions. Secret Agent Society [10] is another commercial serious game for eight to twelve years old HFA children. The game features human and animated characters as well as interactive activities that teach children how to recognise and control their emotions and cope with social challenges such as handling bullying and talking to strangers. Studies with users have shown positive results, both with respect to learning social skills and to generalising the skills learned within the environment. Finally, the Transporters [24] is a three-dimensional animation series created to enhance the understanding and recognition of emotions by children with ASCs between the ages of three and eight. Each of the fifteen five-minute episodes focuses on a different emotion, which is demonstrated through facial expressions by one of the eight characters of the series. The characters are all vehicles with real human faces superimposed on them. After watching the episodes, the child interacts with the system by either matching faces to faces or matching faces to emotions or matching situations to faces. Evaluation studies show that the use of The Transporters led children with HFA to improve significantly in their emotion comprehension and recognition skills for the 15 key emotions presented by the game.

3.3. Serious games for improving interaction skills

Interaction skills such as turn-taking, imitation and collaborative play have been the focus of a number of games for children with autism. The Collaborative Puzzle Game [9], for example, is a tabletop interactive system developed to foster collaboration skills. The game revolves around an interaction rule, called “enforced collaboration”, that establishes that puzzle pieces must be touched and dragged simultaneously by the two players in order to be moved. Results show that the Collaborative Puzzle Game was effective in triggering behaviours associated with coordination of tasks and negotiation in children with autism. Barakova et al. [7] proposed a multi-agent system composed of autonomous interactive blocks that can express emergent behaviours through change in their colours and light intensity as a result of how they are manipulated by the users. These blocks serve as a basis for a variety of educational games that require collaboration between children and focus on explorative behaviour, awareness of others and turn taking. The user tests have shown pronounced explorative behaviour and encouraging results with improvement of turn taking interaction. A number of applications for training interaction skills are not structured games with rules, but simply offer environments and situations that facilitate socialisation. Brigadoon [38], for example, is a virtual island hosted by the commercial online virtual world system known as “Second Life” [43] where individuals with autism can practise their socialisation skills and learn how to interact with each other by manipulating avatars in a safe and risk-free environment. Although a formal evaluation of Brigadoon is not available, anecdotal evidence based on first-hand accounts reported in the Brigadoon forum [38] suggests that Brigadoon is successful among children and young adults with autism because it provides a perceptually immersive socialisation environment, while shielding its users from the anxiety linked with real-world in-person interactions. Parsons et al. [59] developed a virtual environment with two possible scenarios, a virtual cafe and a virtual bus, where users are represented by avatars and can interact with other virtual characters. Studies performed with this virtual environment show that users with ASCs are generally able to use and interpret virtual environments successfully [59], although some children with low verbal IQ and weak executive ability might require more support to complete tasks [60]. In addition, children are guided to learn simple social skills using these two virtual scenarios [48].

Teaching interaction skills to children with autism through games that employ *robotic* toys is an area of emerging interest. A number of studies (e.g., [6,31,73]) have shown that physical embodiment promotes higher social engagement and attribution than virtual embodiment, as well as greater enjoyment, believability and trustfulness, especially with regard to cooperative tasks. Examples of social robots that have been used with children with autism include Keepon [32], a small silicone-made robot. Different robotic platforms have been used in the context of the AuRoRa project [16], from simple mobile robots to more anthropomorphic creatures. All these robots have been used in unstructured play sessions with children with autism and aim to engage them in spontaneous dyadic and triadic interactions, imitation and turn-taking. Usually, the robots were accepted enthusiastically by the children and played the role of social mediators between the children and their caregivers as children often sought to share their excitement towards the robots with the co-present caregivers.

3.4. Agent technology for autism

Although in the last ten years there has been a growing interest in the potential of artificial agents, 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. Little attention has been devoted to autonomy with the exception of the Thinking Head project [47], which focuses on developing a talking head that teaches social skills through its ability of realistically portraying facial expressions, and the virtual peers, Baldi and Timo [14], which are three-dimensional computer-animated talking heads for language and speech training. Furthermore, in line with some behaviourist clinical intervention frameworks, most technologies developed so far 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. This approach does not situate skill acquisition in believable social situations, resulting in a diminished demand for autonomous agents in this context. However, autonomous agents carry a significant potential for autism intervention because they can contribute to the intensive one-on-one support that these children need

while easing the demand for such support from practitioners and parents. Autonomous agents may be able to complement the traditional intervention methods allowing human practitioners to focus on the most complex aspects of face-to-face interventions, while supporting repetitive tasks and on-demand access. On-demand access can be of benefit given abnormal sleep patterns [75] and frequent need for intensive one-on-one support in this population [68,18]. A fully autonomous agent capable of interacting with a child without the presence of an adult could help parents cope with the significant demands of caring for a child with ASCs.

The approach presented in this paper focuses on the development of an *autonomous* agent, i.e. an agent that is able to decide independently how to act best in order to achieve a set of high-level goals that have been delegated to it. Autonomy involves a broad spectrum of behaviours with *no autonomy* and *full autonomy* lying at its extremes. As described in Section 6, the current version of the ECHOES agent falls in between those two extremes: the agent acts autonomously as long as it receives goals to accomplish as well as a description of what actions are available to it and how these actions change the world. In Section 6 we describe the way in which our agent makes independent decisions and acts autonomously. The design of this agent rests on the recommendations from the best autism practice and the SCERTS framework, whereby an optimal interaction style for children with autism “is one that provides enough structure to support a child’s attentional focus, situational understanding, emotional regulation, and positive emotional experience, but that also fosters initiation, spontaneity, flexibility, problem-solving, and self-determination” [63, p. 309]. We argue that these recommendations are in line with the classic agent theory by Wooldridge and Jennings [79], whereby, in addition to autonomy, an agent should be equipped with: (i) *Pro-activeness*, i.e. an ability to exhibit goal-directed behaviour by actively trying to accomplish its goals and taking the initiative; (ii) *Reactivity*, i.e. the ability to perceive the changes in the environment and react to them in a timely manner; and (iii) *Social ability*, i.e. the ability to coordinate its actions with those of another agent – in our case the child. The agent’s social ability is crucial to maximising the chances of the child to experience a sense of self-efficacy in communicating with the agent, including sharing intentions and feelings. Pro-activeness is important to maintaining the child’s attentional focus and to foster motivation, while reactivity is fundamental to adapting the support to the children’s changing needs as well as cognitive and affective states. Therefore, an optimal interaction for children with ASCs can be approximated by an artificial agent equipped with social ability and characterised by the right balance between pro-activeness and reactivity.

4. Pedagogical underpinnings for the ECHOES agent

In order to identify the particular social communication skills that a virtual agent needs to possess to act as a credible social partner to children with ASCs and to support their social competencies, we draw from the relevant knowledge of autism research. Specifically, we build on SCERTS [63], a comprehensive approach to social communication assessment and intervention in autism, to inform the design of the ECHOES agent and to structure the interaction of the agent with the child. The specific questions we ask are: What are the defining characteristics of a social interaction between an agent and a child? What behaviour should the agent exhibit in order to achieve such an interaction?

SCERTS builds on many established intervention methods currently in use and has evolved from a combination of psychological research, clinical practice and educational intervention experience. The framework identifies the particular skills that are essential to successful social communication and that we argue are also necessary for an *ideal* virtual agent to act as a credible social partner to children with ASCs and to help them improve their social competencies. These skills are encapsulated in three overarching domains:

- *Social Communication (SC)*: spontaneous and functional communication, emotional expression, and secure and trusting relationships with children and adults. The intervention objective related to these skills is to help the child increase competence, confidence, and active participation in social activities.
- *Emotional Regulation (ER)*: the ability of the child to maintain a well-regulated emotional state to cope with everyday stress and thus be available for learning and interacting.
- *Transactional Support (TS)*: the development and implementation of tools to help caregivers respond to the child’s needs and interests, modify and adapt the environment and enhance their learning. TS is addressed in three different interaction contexts: interpersonal support, educational support, and family support.

Furthermore, SCERTS divides the assessment of the child’s abilities and required intervention into three developmental stages: *social partner*, where the child communicates through pre-symbolic means, such as gestures and vocalisations; *language partner*, where the child uses early symbolic means to communicate shared meanings (oral language, sign language, picture symbols); and *conversational partner*, where the child exhibits more advanced language abilities (sentences and conversational level discourse).

SCERTS breaks down each domain into a number of essential constitutive components, for each of which it provides a detailed description of the pedagogical objectives to be achieved, the strategies for intervention and the assessment criteria. We build on this operationalisation of social communication for designing the agent’s behaviour and its interaction with the child.

Table 1

Pedagogical goals for the ECHOES learning activities.

Component	Objective
Joint attention	<ul style="list-style-type: none"> • Engage in reciprocal interaction by initiating/responding to bids for interaction • Share attention by looking toward people, shifting gaze between people and objects and following contact/distal point • Share emotions by using facial expressions or vocalisations • Share intentions to regulate behaviour of others by requesting or refusing desired objects or actions • Share intentions for social interaction by greeting, turn taking, and attracting attention • Share intentions for joint attention by commenting on objects or events
Symbol use	<ul style="list-style-type: none"> • Learn by imitation of familiar actions and sounds • Understand non-verbal cues (such as facial expressions, intonation cues, and gestures) in familiar activities • Use gestures and non-verbal means to share intentions (use proximity, facial expressions, contact/distal gestures, coordinate gestures and gaze) • Understand a few familiar words by responding to own name, object names and to a few frequently used phrases in familiar routines

5. The ECHOES learning activities

The interaction between the child and the agent is structured around a number of different *learning activities* that are facilitated by a 42 inch multitouch LCD display with eye-gaze tracking. For each child, the learning activities are selected manually by a human operator (practitioner, parent or other carer) through a graphical interface. The ECHOES learning activities focus on social communication and, in particular, on the two sub-components of social communication that have been identified by SCERTS as the most challenging for children with ASCs:

- (i) *Joint attention*: Child's ability to coordinate and share attention by looking towards people or shifting gaze between people and objects, share emotions by using facial expressions, express intentions, engage in turn-taking and participate in reciprocal social interactions by initiating/responding to bids for interaction; and
- (ii) *Symbol use*: Child's understanding of meaning expressed through conventional gestures, words, and more advanced linguistic forms; child's ability to conventionally use objects in play; and child's ability to use non-verbal means and vocalisations to share intentions.

ECHOES includes a total of twelve learning activities designed to enhance the child's capacities for joint attentional interactions and symbol use and target one or more of the pedagogical goals listed in [Table 1](#).

These goals directly correspond to the intervention goals specified in the SCERTS framework, although we adapted some of the goals in order to fit the human–computer interaction context of our system (SCERTS has been developed for a human–human intervention context instead). Given that our target population is comprised of young children aged between five and seven, the ECHOES learning activities focus on the SCERTS objectives that pertain to the social and early language partner stages, and deliberately omit objectives relating to the conversational partner stage.

All the ECHOES activities take place in a sensory garden populated by Andy and by interactive “magic” objects that react in unusual ways sometimes transforming into other objects when the agent or the child act upon them through specific touch gestures (see [Fig. 1](#)). For example, tapping the petals of a flower makes the flower become a floating bubble or a bouncy ball. The relationship of action–reaction implemented through touch promotes children's understanding of cause and effect. ECHOES recognises the following touch actions: single touch, touch with multiple fingers, one or two finger dragging, and holding an object on top of another. Although we do not represent the world literally in ECHOES, we maintain the consistency of the behaviours of the different objects within the environment in order to support children's understanding of cause and effect. The decision to set the learning activities in a magic garden serves a number of purposes. Through the mystery of the enchanted garden, ECHOES motivates the children to interact with the environment and to engage with the tasks. Since it is not obvious to the children from the outset how the different objects will behave and react to their gestures, the sensory garden fosters children's imaginative play and exploratory behaviour, which are rarely observed in individuals with ASCs [77]. In this respect, the ECHOES environment is unique in the context of technology-enhanced intervention for autism, where the onus of intervention tends to be placed mostly on training through repetition of specific stimuli (e.g. [24]). 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 [64]. We reinforced this principle by developing activities that are linked in non-verbal *narratives*. For example, in one activity flowers are transformed into bouncy balls, in another activity the colour of bouncy balls can be changed by having them thrown through a cloud, and in a third activity different coloured balls are sorted into containers. The activities are distinct from each other and call for an active participation of the child who plays a central role in choosing further activities.

As for the specific contents of the learning activities, we follow the SCERTS philosophy advocating that learning activities need to be “meaningful and purposeful” [64], in contrast with approaches in which the activities are task-based and skills are trained in a repetitive fashion and in isolation from a meaningful context. We designed two sets of activities:




Fig. 1. Children interacting with Andy through the ECHOES multi-touch display.

- *Goal-oriented* activities, with clear sequence of steps and an easily identifiable end-goal; and
- *Cooperative turn-taking* activities, with no clear end-goal and whose main objectives are social reciprocity, turn taking, and mutual enjoyment.

Sorting a set of balls according to their colours and collecting all the flowers on the ground into a basket are examples of goal-oriented activities, while taking turns with the agent in order to grow flowers by shaking a cloud that produces rain and throwing balls thorough a cloud so that they change colour constitute turn-taking activities. All activities are supposed to be performed by Andy and the child in cooperation, with Andy assuming a more or less prominent role according to the particular learning objective of an activity. For example, if the goal is *respond-to-reciprocal-interaction*, Andy will adopt a leading role in prompting the child to take a turn. If the goal is *initiate-reciprocal-interaction*, Andy will wait to give the child an opportunity to initiate a bid for interaction, before initiating the interaction if the child is not interacting. In the next section, we shall describe the full repertoire of Andy's behaviours.

The ideas of a magic garden and of morphing objects as well as the content of the specific activities have emerged from participatory design workshops with both typically developing children and children with ASCs and knowledge elicitation workshops with autism experts and practitioners. More specifically, fourteen participatory design workshops involving eighty-seven typically developing children and fifty-three children with ASCs were conducted during the lifespan of the project with the goal of informing the design of the look-and-feel of the environment including the functionality and the interactive properties of objects, the appearance of the ECHOES' agents and other aesthetic decisions [22]. In addition, we conducted two knowledge elicitation workshops involving thirty practitioners with extensive experience of day-to-day autism intervention and in-depth knowledge of a large number of individual children with ASCs and three older (eleven-eighteen years old) high functioning children and teenagers with ASCs who acted as consultants. These workshops focused on the appearance of the agent, its specific behaviours toward the children and the narratives underlying the learning activities. Finally, the ECHOES environment and a subset of the learning activities were tested at different stages of the project in four formative evaluation studies involving forty-six children with ASCs [2]. A number of crucial design decisions, including the appearance of the agent, have been influenced by these formative evaluation studies.

6. The design and implementation of the ECHOES agent

Our goal was to create an artificial social partner that could act credibly both as a peer and as a tutor to children with ASCs and, as a result of this, that could deliver the educational and interpersonal support advocated in SCERTS. As a *tutor*, our agent aims to deliver visual and organisational support for:

- Expanding and enhancing the development of a child's expressive communication system" [63, p. 309];
- Supporting a child's understanding of language as well as others' non-verbal behaviour" [63, p. 309]; and
- Supporting a child's sense of organisation, activity structure, and sense of time" [63, p. 309].

In addition, the ECHOES agent acts as a *peer* to the child and aims to provide them with interpersonal support by:

- Accommodating the child's preference for structure and predictability, while fostering initiation, spontaneity, and self-determination; and
- Exposing the child to a positive interaction with a peer so that they can "benefit optimally from good language, social, and play models" [63, p. 309].

To achieve this goal, we designed the agent's physical appearance, its behaviour and its interaction with children based on the SCERTS model and input received from thirty ASCs practitioners, who participated in two workshops

involving storyboarding tools and group discussions, as well as individual interviews organised over the lifespan of the ECHOES project. In what follows, we will describe the main design choices that we made to create Andy. Additional technical information on the agent's architecture can be found in Bernardini et al. [13] and Bernardini and Porayska-Pomsta [11].

6.1. Agent's intelligence

Among the various domain-independent agent architectures that have been proposed for building agents, FATiMA [20] is particularly well suited to fulfil the design requirements of our agent, because it combines the kind of *reactive* and *cognitive* capabilities needed to implement an autonomous, proactive and reactive *socio-emotionally* competent agent, like the one we had envisaged for this context. The cognitive layer of FATiMA is based on artificial intelligence *planning* techniques, in particular an implementation of the classical partial order causal link planning algorithm [69], while the emotional model is derived from the OCC theory of emotions [53] and the appraisal theory [70]. A FATiMA agent is characterised by:

- (1) a set of internal goals;
- (2) a set of action strategies to achieve these goals; and
- (3) an affective system which includes:
 - (i) emotional reaction rules, to determine how generic events are appraised by the agent;
 - (ii) action tendencies, to represent the agent's impulsive actions to different emotional states;
 - (iii) emotional thresholds, to specify the agent's resistance to different emotions; and
 - (iv) decay rates for emotions, to determine how fast an emotion decays over time.

The two main mechanisms controlling a FATiMA agent are *appraisal* and *coping*. The agent experiences one or more of the 22 emotions of the OCC model [53], based on its appraisal of the current external events with respect to its own goals and based on 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 two different levels: the *reactive* level, which affects the short-term horizon of the agent's behaviour, and the *deliberative* level, which relates to the agent's long term goal-oriented behaviour.

In ECHOES, each learning activity has a FATiMA agent model associated with it. All these models share the same specification of the agent's affective system, because we want the agent to maintain the same personality from session to session in order to establish a trusting relationship with the child. Andy is a positive, motivating and supportive character, which we obtained by manipulating its affective system as well as its goals and action strategies. For example, Andy always greets the child by name when it walks in, gives positive feedback to the child when the child initiates or responds to a bid for interaction, and tries to reengage the child if the child gets distracted. Andy tends to be happy and does not get frustrated easily, therefore its happiness threshold is low, while its frustration threshold is high, and its happiness decay is slow while its frustration decay is fast. We control Andy's facial expressions and gestures through the specification of Andy's action tendencies. For example, the agent smiles and gives a thumb up when the child engages in an activity with it.

While Andy's personality does not change between activities, the set of goals that the ECHOES agent actively tries to pursue as well as its action strategies are specified for each learning activity on the basis of: (i) the high-level pedagogical goals on which the activity focuses; and (ii) the specific narrative content of the activity itself. Andy has separate goals for engaging with the child to support the pedagogical goals and for acting within the context of the activity. For example, if the high-level goal of an activity is "Engage in reciprocal interaction" and the content of the activity involves picking flowers from the garden, one of the low-level goals of the agent will be to fill a basket with flowers in collaboration with the child, whereas its action strategies will demonstrate to the child different ways of engaging in reciprocal interaction, e.g. the agent can choose between pointing at a flower, looking at it, or saying "Your turn!".

The process of authoring Andy through FATiMA was a complex task. The FATiMA characters' behaviour emerges from a combination of multiple factors and, for the final behaviour to be as one intended, not only do operators and goals need to be carefully defined, but *ad hoc* values also need to be assigned to emotional reactions, action tendencies, emotional thresholds and emotion decay rates. Since there are no conventional definitions of such values and no clear rules of how such values interact with one another, they have to be adjusted through trial and error. This process of fine-tuning the domain model is time-consuming and, often, cumbersome. This is not a problem specific to ECHOES and other teams using FATiMA for authoring characters reported similar problems (e.g. [33,34]).

6.2. The ECHOES' intelligence engine

Although the FATiMA architecture represents the fundamental base of the ECHOES agent's autonomous behaviour, the agent exhibits a higher level of intelligence when FATiMA is complemented by the two additional components that, together with it, constitute the ECHOES' *intelligent engine*:

- the *pedagogic component*, which is intended to monitor the unfolding of the interaction between the child and the agent to ensure the achievement of the learning objectives; and

- the *child model*, a user model that aims to assess the cognitive and emotional state of the child in real time in order to constantly feedback this information to both the agent and the pedagogic component.

We first describe the design of these components in detail, and then report the challenges that we faced in implementing these components. Additional information concerning the ECHOES intelligent engine can be found in [12].

One of the roles of the pedagogic component is to establish the initial state and the goals for each session and each user. While the overall set of goals within which the pedagogic component can choose is formulated on the basis of the SCERTS framework, both the initial situation and the specific set of goals for each user in a given session need to be decided on the basis of the user's profile and their interaction history with the ECHOES system. After delegating the session goals to the agent, the pedagogic component is required to leave the agent free to interact with the child without interfering. However, it externally monitors the unfolding of the events and, at the same time, receives input from the user model about the child's mental state. If the interaction between the child and the agent significantly diverges from the achievement of the pedagogical goals for that session or the child is observed to experience highly negative mental states such as extreme anxiety or arousal, the pedagogic component is allowed to intervene to keep the interaction on track. It can, for example, influence how the planner constructs the plan, change the overall goals of the session and even drop these goals, if appropriate. Conceptually, the ECHOES' pedagogic component acts similarly to the "drama manager" in interactive storytelling systems [67], although its role is more limited. In such systems, the drama manager monitors the unfolding of the story that emerges from the interaction between the user and the characters and can actively intervene to preserve the coherence of the story by changing the behaviours of the characters.

The role of the child model is to estimate the cognitive and affective state of the child and feed this information to the agent and the pedagogic component whenever a change is detected. A user model is key for realising the SCERTS' principle whereby an "optimal style of interaction" focuses on "supporting children to be as successful as possible in experiencing a sense of efficacy in communicating their intentions, and in participating in affectively charged and emotionally fulfilling social engagement with a variety of partners" [63, p. 309]. This is because the child model supports Andy in making informed decisions about what to do next based on the current mental state of the child and in reacting according to its estimation of the child's intentions, needs and desires. In ECHOES, the child model aims to infer the mental states of the child based on real-time information coming from the touch and eye-tracking systems and produces output at two levels: cognitive and affective. The cognitive assessment is facilitated by a rule-based engine which estimates the extent to which the child has achieved the pedagogical goals associated with the session. The rules are based on the SCERTS model, which provides clear guidelines and precise timing constraints to establish whether the child had mastered specific skills in relation to joint attention and symbolic use. The affective assessment is facilitated by a combination of supervised and unsupervised learning techniques used to estimate the child's level of engagement with the system among five different categories, from complete engagement to total disengagement. Engagement is an important indicator of the affective state of a child because disengagement is usually linked with boredom and anxiety, whereas engagement with interest and excitement. In order to obtain data to train the classifier, we conducted two studies with children with autism using a first prototype of ECHOES and involving a total of 46 children aged 5 to 14. The interactions between the children and the environment were video recorded and then annotated for engagement. We synchronised the video annotations with the system's log files, concentrating particularly on how often the child touched the screen. Using these data, we then trained a Support Vector Machine (SVM) classifier using Weka to predict engagement. Every second, the classifier estimates the participant's level of engagement based on how often the participant touched the screen in the preceding one to five seconds. Our tests of engagement estimation based on unsampled data from six children interacting with ECHOES suggest 68% accuracy for a baseline classifier for an overall F-measure of 0.078 using 10-fold cross-validation (see [62] for more details on the child model).

We faced significant challenges in implementing the child model, particularly in interpreting the child's mental state accurately. In general, reliable interpretation of complex human mental states at a deep level is a complex and still open problem [56,55,54], and little effort has been devoted to modelling young children cognitive and affective states (see [80] for one of the few examples of user modelling targeting children). In ECHOES, the challenge of assessing young children mental states was made even more difficult by the naturalistic context in which the system was used: with the child standing and being allowed to move freely, it was difficult to collect reliable data through the touch and eye-tracking systems. Unreliable output from the child model has the consequence of diminishing the ability of the pedagogic component and the agent to perform their respective tasks timely and appropriately, as their actions are based on feedback information concerning the child's current mental state. In order to cope with the child model failures, we supplemented the system with a wizard-of-oz control panel through which a human operator has the opportunity to alter the behaviour of the agent without relying on the pedagogic component's intervention. Through the control panel, the operator can also establish the initial state and the goals for the session and change activity when needed. We used this version of the system for the evaluation studies reported in Section 7.

6.3. Agent's physical appearance

The physical appearance of Andy was established based on a combination of research studies (see e.g. [2]) and participatory design workshops with six teachers in one of the UK autism special primary schools. The practitioners were presented with a prototype of the ECHOES system and an early version of the agent, called Paul (see Fig. 2).

They were first asked to interact with the system to get the feel for what an interaction with the ECHOES hardware and software entailed. Following this interactive session, the practitioners were divided into pairs and were asked to storyboard

the possible activities for ECHOES (see Fig. 3) using real children's profiles, which were prepared by them *a priori* and anonymised.

To aid the storyboarding activities and to obtain as detailed information as possible concerning the behaviours needed for the agent, the teachers were asked to role-play, with one practitioner playing the role of the child and the other playing the role of the agent. Both teachers were given a printout of the background of the ECHOES magic garden with cut-outs of the different ECHOES objects and of Paul displaying different gestures and facial expressions (see Fig. 2). In addition, the practitioners who were enacting the agent were given its profile, which outlined what the agent could and could not see, say and do.

The teachers could then stick the different cut-outs on top of the ECHOES background in the way they felt appropriate, thus constructing both the activities and the specific actions of the agent. We asked them to critique the flexibility of the designs in terms of both the environment and the agent. The designs and the critique were captured through videos, photographs, researcher notes and finally the storyboards themselves, which, crucially in respect to the agent's appearance, contained annotations on top of the cut outs suggesting the improvements needed for the agent.

Twelve different storyboards were constructed by the six teachers participating in the workshop, each based on a different child's profile. In relation to the agent, the outcomes were consistent between the storyboards. The practitioners recommended that the ECHOES agent should have a *child-like* physical appearance. An agent with such an appearance carries the advantage of a predictable child-like companion and provides the children with the possibility to create a history of successful interactions between them and a peer. This may provide the basis for motivating the child to relate to other children in real life. These recommendations are in line with current research and practices. In particular, children with ASCs often avoid interactions with other children who are usually less predictable than adults. SCERTS emphasises the importance of creating opportunities for children with ASCs to play with other children and to develop positive relationships with them [63].

Furthermore, the ECHOES agent was given a *cartoonish* look, which was deemed by the teachers as potentially more familiar and more fun to the children than a photo-realistic alternative, thanks to children's familiarity with television and computer games characters. This is also in line with recent studies reporting that individuals with ASCs and their care-givers recommend that assistive software should be designed with fun in mind [65]. The teachers also agreed that the agent's head and eyes should be bigger than the rest of its body and that the agent should have the ability to move its eyes in different directions in order to allow the agent to provide the children with unambiguous non-verbal social cues, e.g. through looking at a particular object. Such cues tend to be missed by children with ASCs in real life. Finally, the agent's hands needed to be detailed enough both in appearance and movement to facilitate unambiguous pointing. These recommendations resulted in a new ECHOES agent, called Andy, as shown in Fig. 4. The model for Andy was sourced from TurboSquid [46].

6.4. Repertoire of the agent's behaviours

Given that the focus of ECHOES is on supporting children's social communication, in particular, joint attention and symbol use, the agent's actions are either concrete demonstrations of the related skills or actions performed to invite the child to practice those skills. Specifically, we define the joint attention and symbolic use in terms of three component skills:

- (i) Responding to bids for interaction;
- (ii) Initiating bids for interaction; and
- (iii) Engaging in turn taking.



Fig. 2. First prototype of the ECHOES environment and the ECHOES agent, Paul, used in participatory design workshops with practitioners.



Fig. 3. ECHOES storyboarding activities with practitioners.



Fig. 4. Andy, the final appearance of the ECHOES virtual agent.

Our agent is able to perform these skills in three different ways:

- (a) Verbally by using simple language or key phrases (e.g., “My turn!” and “Your turn!” for turn-taking);
- (b) Non-verbally through gaze and gestures such as pointing at an object from a distance or touching an object; and
- (c) By combining verbal and non-verbal behaviours.

Initiating a non-verbal bid for interaction by the agent involves Andy looking at the child, then looking at an object followed by indicating that object and finally looking back at the child. The ECHOES agent is able to make requests, to greet the child by name, to comment on actions or events happening in the garden and to use exploratory actions on objects, e.g. to explore the features of the magic objects populating the garden. This variety of behaviours makes the interaction dynamic enough to keep the child engaged and, potentially, to foster generalisation, while retaining a degree of predictability that is essential to supporting the child’s attentional focus.

The practitioners emphasised the importance of providing the children with *positive feedback* in order to reduce their anxiety related to social interactions and to help them experience a sense of self-efficacy and achievement. Andy always provides the child with positive feedback, especially if the child follows the agent’s bids for interaction correctly in task-based activities. If the child does not perform the required action, the agent first waits for the child to do things at their own pace and then intervenes by demonstrating the action and encouraging the child to try again. To provide organisational support, the agent always explains a new activity to the child by using a simple language and precise instructions (e.g., “Let’s pick ten flowers”).

Finally, Andy is a *responsive* agent, which is regarded as crucial by practitioners. Its responsiveness ranges from simple physical reactions to actions performed on its body (for example, Andy laughs if the child tickles it), to the agent’s ability to respond to the child’s needs (for example, Andy repeats a bid for interaction if the child does not respond to it within a sensible amount of time). Ideally, the agent should attune its emotional tone to that of the child in order to keep emotional

engagement with the child and assist them during emotionally arousing experiences. Such a sophisticated level of reactivity requires that the agent is able to assess in real-time the current cognitive and emotional state of the user at a fine-level of granularity. As described in Section 6.2, ECHOES includes a *user model*, but this model is currently very simple and, therefore, is not able to assess subtle and nuanced emotional states.

6.5. Verbal and non-verbal communication

Our agent uses (a fragment of) the *Makaton* language [74] to facilitate communication with the child. Makaton is a language programme of signs and symbols to support spoken language. Makaton is often used with individuals who find it difficult to communicate verbally and is extensively employed in special education schools across the UK, but also in other countries worldwide. Some signs are straightforward, for example “Yes” is accompanied by a head nod and “Good job!” by a thumb up, while other signs are more complex, for example “Your turn” is indicated by the hand held in fist and by the base of the hand pointing toward the person being addressed. In general, *hands* are powerful communication mechanisms and the ECHOES agent uses them to augment its verbal language graphically.

Auditory information can be challenging for children with ASCs. Therefore, we kept the *spoken language* as simple as possible. We also sought help from a small group of typically developing children in order to make the language sound natural to other children. We use “My turn!” and “Your turn!” for facilitating turn-taking, “Good job!”, “Wow!”, “Cool!” for showing enthusiasm, “All done!” to signal the end of an activity. All bids for interactions are extremely simple, for example “Touch the cloud” and “Let’s grow flowers in these pots”.

The agent can perform a number of positive *facial expressions*: it can smile, laugh, and look happy. We also have a neutral expression used when a positive expression would not be appropriate. These expressions are implemented by careful changes in the agent’s lips and eyebrows. They are usually accompanied by *body gestures* consistent with emotion showed through the face in order to reinforce the message. For example, the agent smiles when it gives positive feedback to the child through a thumbs-up gesture. In order to make facial expressions and the gestures more evident to children, they are all slightly exaggerated compared to natural human behaviours. However, given that over-sensitivity and distress from extreme sensory stimuli is common in ASC [30], none of Andy’s emotional displays are complex.

7. Experimental results

We present experimental results pertaining to the effectiveness of Andy in acting as a social partner to children with autism based on an extensive evaluation of the ECHOES platform, which show encouraging tendencies for a number of children.

7.1. Experimental design

To assess the impact of Andy and the ECHOES environment on social communication skills in children with ASCs, a large scale multi-site intervention study was conducted. The system was deployed to five special schools and schools with units dedicated to working with children with ASCs across the UK. The decision was made to not attempt a full-scale randomly controlled trial (RCT) of ECHOES due to the difficulty in matching individuals across the groups given the within-ASC heterogeneity. At this stage, we are more interested in how children might use ECHOES and whether it would elicit any social communicative behaviours not observed outside of ECHOES.

In order to investigate these factors, ECHOES was installed in quiet, dedicated spaces within each school in which individual children could interact with the environment whilst their interaction was monitored and structured by a human practitioner. To assess each child’s initial social communication skills, the child’s behaviour in several natural and semi-natural environments was documented. Each child was videoed during: (i) free-play (e.g. in the playground); (ii) free-activity in classroom; (iii) a structured group turn-taking exercise in the classroom; and (iv) a structured one-on-one table-top turn-taking activity. Behaviours observed in these videos were coded and quantified (see Section 7.2). Due to the time and complexity involved in video annotation, only the results of the one-on-one table-top activity will be presented here. The structured table-top turn-taking activity involved the child playing with two toys (a bubble gun and a remote-controlled robot) on a table-top with a human practitioner (see Fig. 5). The practitioner was instructed to take turns with the child in controlling the toy, using joint attention and pointing to attempt to direct the child’s attention to objects on the table-top, and respond to any bids for interaction made by the child.

After the initial table-top pre-test, each child was given the opportunity to play with the environment for periods of 10–20 min, several times a week over a six week period. The child would sit or stand in front of the ECHOES screen whilst the practitioner sat to the side of the screen out of their immediate line of sight (see Fig. 6). The structure of each session, in particular the decision of which learning activities to engage with and when to progress between them, was decided collaboratively by the child and the practitioner. Andy was present during most of the learning activities and varied in the extent it was critical to the goal of an activity. For instance, during the early exploratory activity, Andy would be present and the child could tickle or hand it objects, but it would not initiate interaction with the child. Later activities such as sorting coloured balls into pots required turn-taking and Andy would request balls which the child could hand to it.

If the child became distressed or disengaged with ECHOES during the sessions, the session was terminated and the child returned to their class. If this happened frequently and impeded the child's use of ECHOES, the child was removed from the final analysis sample.

ECHOES sessions were videoed. At the end of the ECHOES six week usage, a second table-top session was conducted to assess generalisation of the social behaviours learned during use of ECHOES. Analysis was based on the pre-ECHOES table-top video, three 15 min ECHOES sessions involving Andy (beginning, middle, and end of the intervention period), and the post-ECHOES table-top.

7.2. Video annotation

In order to describe each child's broad socio-communicative abilities a coding scheme was required that could be applied to both the pre/post-table-top videos and the ECHOES interactions. Existing diagnostic frameworks such as ADOS require specific testing protocols, do not provide the fine-grained assessment required to monitor change and are focussed on diagnosis which was not the objective of this study. Instead, the SCERTS Assessment Protocol (SAP) was used. SAP is a curriculum-based assessment designed to measure a range of behavioural and emotional outcomes in contexts that are meaningful and functional, such as classroom activities, free time on the playground, or interacting with family at home. SCERTS assessment focusses on a child's achievements, rather than deficits, in particular domains such as degree of success in communicative exchange, related dimensions of emotional expression and regulation, and the child's social competence and active participation in natural activities and environments. Assessment of a child's abilities within the SCERTS framework is usually conducted using the functional role of various behaviours and communicative acts are emphasised, instead of specific milestones of language/skill acquisition, as well as consistency with which behaviours are demonstrated across contexts and with different social partners.

The existing SAP protocol scores ability on a 0–2 scale: 0 = Behavioural criteria is not met, 1 = Behaviour criteria is met inconsistently, 2 = Behaviour criteria is met consistently. After initial piloting of this scoring system with the table-top and ECHOES videos, it was decided that such a coarse score did not capture the nuance of the child's interactions with Andy or the human practitioner. Coding an entire interaction in terms of whether the child manifested a behaviour or not also restricted how much we could learn about the conditions under which the child performed that behaviour. As such, the SAP protocol was extended significantly. All existing behavioural categories (such as "Child responds to verbal and non-verbal bids for interaction") were retained whilst new behaviours of interest to ECHOES were introduced (such as "Child uses a pointing gesture"). The revised coding scheme was formalised in a comprehensive coding manual containing definitions of all behaviours and how they should be applied. This coding scheme will be discussed in depth and published elsewhere. Coding was performed using the software "ELAN" (professional audio and video annotation tool). Each video was assessed for sixteen main behavioural categories. When an instance of a behaviour was identified, the appropriate code was associated to the period of the video containing the behaviour. This coding scheme allows patterns of behaviours and their frequencies to be identified, such as social interactions initiated by the child or instances of turn taking.

First coding was performed by ten independent coders who had completed a training period with the adapted SAP. Each video was moderated by a second coder who assessed the validity of each code and suggested deletions or additions. Disagreements were resolved through discussion with a third coder. Twenty percent of all videos are currently undergoing blind second coding and inter-rater reliability will be reported in future reports.

7.3. Participants

Twenty-nine children were recruited from five special units in mainstream primary schools or schools dedicated to providing care for children with ASCs and/or other disabilities across the UK. Each child had previously received a community clinical diagnosis of autism spectrum disorder or aspergers before being referred to the special educational unit. For the analysis presented here, a subset of nineteen ASC children who received the most exposure to ECHOES and completed




Fig. 5. The table-top activity used to assess social communication skill before (pre) and after (post) ECHOES use. Left = bubbles activity. Right = robot activity.



Fig. 6. Children using ECHOES with a practitioner. Children sat on a seat in front of the screen whilst the practitioner (and classroom assistant when necessary) sat to the side of the screen out of the immediate line of sight of the child. Both images depict moments in which the child spontaneously directed a statement to the practitioner (“I’ve got no shoes today”, left image and “Andy is going home now”, right image).

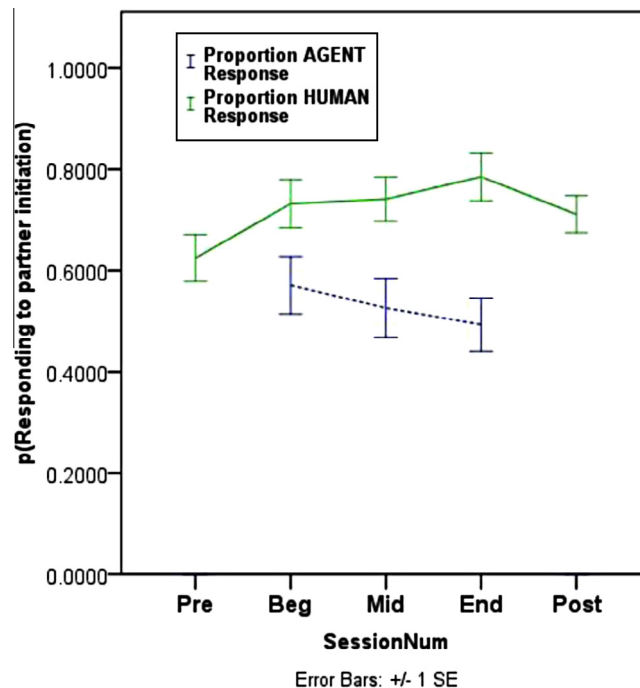


Fig. 7. Probability that child responds to an initiation by the human practitioner (solid green line) or agent (dotted blue line) across the pre table-top, beginning, middle and end ECHOES sessions and the post table-top session. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the pre and post table-top assessments were selected. The majority of children dropped out due to illness and absence although a few refused to continue to use ECHOES. These children were not distinguished from the remaining children by any diagnostic traits.

The analysis group had an average chronological age of 8 years 5 months (range 4 to 14 yrs) and contained one girl. Confirmation of the ASC diagnosis was provided by the Social Communication Questionnaire (SCQ) completed by a caregiver. Average SCQ score was 23.4 (SD = 4.64; Range = 14–32), higher than the 15 cut-off for possible ASC. Verbal language ability assessed using the British Picture Vocabulary Scale (BPVS) reported a group mean of 36.62 (Raw) and age equivalent of 3.99 (SD = 0.97), significantly lower than their group age. It should be noted that seven out of nineteen children scored too low on the BPVS to provide an age equivalence.

7.4. Results

As the children varied in their number of ECHOES sessions 15 min periods during which they interacted with Andy were identified from the beginning, middle and end of the six week intervention period. Video annotations were applied to each 15 min video using “ELAN” and moderated by a second coder.

Among the sixteen main behavioural categories contained by the modified SAP coding scheme, due to space limitations, we will focus on only two main social behaviours, which are often severely impaired in children with ASCs:

- (i) Responding to Andy's bids for interaction (responses); and
- (ii) Bids for Andy to interact (initiations).

7.4.1. Child's response to social partner

In order to identify whether children were responding to partner bids for interaction, pairs of clear partner initiations and child appropriate responses were identified in the video codes. Each session varied in length and content within and across children creating uninformative changes in the frequency of partner initiations. To control for this child responses were calculated as a probability of responding appropriately to each initiation across a whole session: a high probability indicates more reliable response to initiations, a low probability indicates more failure to respond. Child actions that were unrelated to the partner initiation, a continuation of an existing action or off-task (e.g. requesting the toilet) were not coded as responses.

The mean probability that the child responded to the practitioner's bids for interaction during the table-top pre-test was 0.66 (SD = 0.17). After the intervention the probability was 0.71 (SD = 0.14) (see Fig. 7). This slight increase in responses between the pre and post table-top test was not significant; $t(14) = -1.637, p = .124$, n.s. (note: four participants were excluded from this analysis due to missing data). However, during ECHOES the probability that the child responded to the practitioners bids for interaction increased significantly relative to the table-top pre-test: beginning = 0.73 (SD = 0.21, $t(16) = -3.107, p < .01$), middle = 0.74 (SD = 0.19, $t(16) = -3.387, p < .01$), end = 0.79 (SD = 0.21, $t(16) = -4.072, p < .001$). The slight increase in response probability across the three ECHOES sessions (+0.06) was non-significant (all $t_s < 1$). This increase relative to the pre table-top may suggest a comfort with the ECHOES environment that elicited a level of responsiveness in the child not observed during the table-top activity. This comfort cannot be due to a simple familiarity with the practitioner as the practitioner would often change from session to session. The role Andy played in facilitating these interactions with the human practitioner is unclear. Across the three ECHOES sessions, the proportion of Andy's initiations responded to by the child showed a slight but non-significant decrease: beginning = 0.57 (SD = 0.22), middle = 0.53 (SD 0.25), end = 0.49 (SD 0.23); all $t_s < 1$. Given the increasing complexity of the learning activities and their dependence on turn-taking with Andy across the three sessions, a growing disinterest in Andy could have been predicted (possibly reflected by the numerical decrease in response probability) but the relative stability of the child's response rate to agent initiations is reassuring. We are exploring further refinement of Andy's intelligence and diversity of behaviour to investigate whether this downturn in response rate can be avoided.

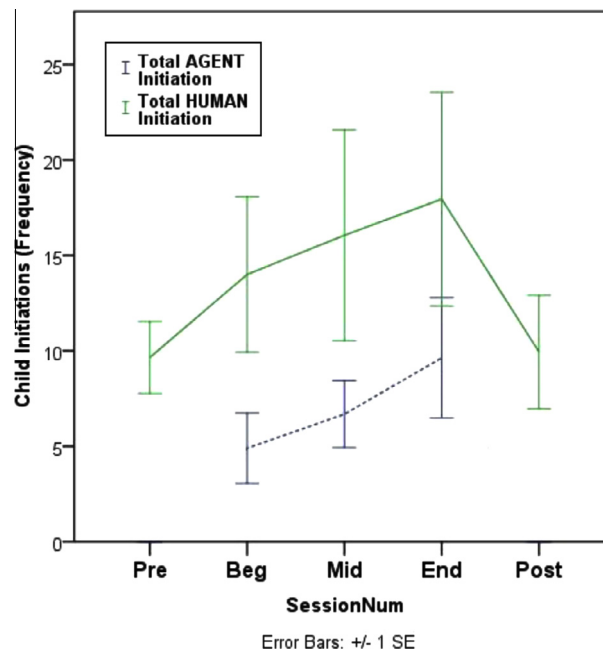


Fig. 8. Frequency of child initiations to the human practitioner (solid green line) or agent (dotted blue line) across the pre table-top, beginning, middle and end ECHOES sessions and the post table-top session. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



Fig. 9. A child, who was believed to be non-communicative, waved and said “Hi Andy!” when the agent walked on the screen during one of his last sessions with ECHOES.

7.4.2. Child's initiations to social partner

Child responses place most of the impetus on Andy or the human practitioner. A more advanced social behaviour rarely observed in children with ASCs is the initiation of social interactions. As child initiations are not contingent on the partner's behaviour this measure can be expressed as the raw frequency of initiations coded in the video for each session. The frequency with which the child initiated an interaction during the table-top pre-test was low, 9.65 times (SD 7.78) and did not change by the post table-top session, 9.93 times (SD 11.5): $t(18) = .154, p = .88$, n.s. (see Fig. 8).

The number of initiations to the human practitioner numerically increased across the three ECHOES sessions from 14.06 (SD = 18.3) to 16.29 (SD = 25.5) and 17.94 (SD = 25.8) but the large variance across individual children rendered this increase non-significant. Initiations to Andy also numerically increased across the three ECHOES sessions from 4.89 (SD 8.05) to 6.68 (SD 7.68) and 9.63 (SD 13.74). Unfortunately, this difference did not reach significance (beg-end: $t(18) = -1.719, p = .103$, n.s.). Eight children increased their number of initiations to Andy, seven produced the same number and only four decreased. This suggests that the heterogeneity of the children with ASCs who participate in our study may make it difficult to identify a group increase in initiations. However, an encouraging indication of the child's increasing social engagement with Andy can be seen in the decreasing difference in the frequency of initiations the child made to Andy compared to the human practitioner (Fig. 8). During the beginning and middle ECHOES sessions the mean number of initiations made to Andy is significantly lower than the number made to the human practitioner (beg: $t(18) = 2.215, p < .05$; mid: $t(18) = 1.943, p = .068$, i.e. marginal), but by the final ECHOES session this difference had disappeared ($t(18) = 1.621, p = .122$, ns.). This increase might indicate that the child was increasingly regarding Andy as a socially authentic partner who they could interact with. Anecdotal evidence supporting this hypothesis comes from one child who showed no initial interest in Andy, but spontaneously waved and said “Hi Andy!” when the agent walked on the screen in a later session (see Fig. 9). Such behaviours were very surprising to teachers and support workers within the school who believed the child in question to be non-communicative. Similar other cases were observed and reported in relation to other children in the different schools.

Caution should be used when extrapolating broader conclusions for interventions in ASC from our results. Our analysis group contains a relatively small (19) and heterogeneous group (SCQ range = 14–32) of children with autism. The current data is not presented with an IQ or age matched typical control group and, due to the exploratory nature of the ECHOES environment, we could not tightly control the experience of each child across the sessions. As such, the potential for these results to generalise to other children is unknown. However, ECHOES was specifically designed as a flexible and exploratory environment in which children with a range of abilities could explore socio-communication skills. As such, we believe that our results demonstrate that ECHOES and its intelligent agent may be beneficial to individual children.

8. Conclusions and future work

In this paper, we have presented our approach to designing and implementing a serious game for supporting social communication skills of children with ASCs. This approach relies on the use of an autonomous pedagogical agent and best pedagogic practice as implemented in the traditional intervention contexts. The design of the agent and of the ECHOES learning activities is based on principled intervention guidelines, recommendations from autism practitioners and children themselves. Designing a social partner for children, especially children with ASCs, increases the need for credibility and believability of the agent as a social partner. We argued that such believability comes from the ability of the agent to adapt to individual children in real-time and to act autonomously, varying the degree of predictability in order to foster learning and generalisation of the skills acquired in the virtual world to the real environments.

To the best of our knowledge, the ECHOES evaluation represents one of the first major evaluations of a serious game for autism conducted in real-school contexts, across different schools and involving a significant number of children. Although presently we can only report the preliminary coarse-grained analysis of children's behaviours in relation to Andy and whilst

no significant transfer of increased social responsiveness or initiations to real-world contexts were observed across all children, there is evidence of some children having benefited from their exposure to Andy and the ECHOES environment as a whole. Specifically, the experimental results show that the number of initiations that the children make to Andy at the start of using ECHOES is significantly less than the number of initiations that they make to the practitioner, but this difference disappears by the final session. A possible interpretation of this phenomenon is that Andy's reciprocal interactions with the children and his critical role within the learning activities are responsible for eliciting spontaneous social behaviours in the children. Post-intervention interviews conducted with teachers and support workers within the schools that participated in the evaluation, all of whom were enthusiastic about the ECHOES environment, indicate that this is a plausible interpretation. The teachers agreed that observing children interacting with ECHOES and Andy provided them with an unusual and unexpected window onto the individual children's capabilities. One teacher's comment emblematically expresses this impression and summarises many other comments made by different teachers:

We found things through this almost diagnostic tool, as it became, that the children were doing and we never realised that they had these skills, because some of them were so locked in. The normal curriculum that we were offering them, slightly modified obviously, was not allowing them to demonstrate the skills to us, etc. The children with autism would play in a very relaxed state, manipulating Andy in a way that would make them laugh and was perhaps a little obtuse, perhaps a little naughty, perhaps almost the opposite of what they were expected to do. And the fact that they're able to sit and play without any judgement allowed them to have opportunities which they never had before.¹

Furthermore, during the course of the experiments with ECHOES, a number of teachers carefully observed individual children's behaviours with Andy, annotated those observations and then compared them with the typical classroom conduct of the same children. Such comparisons often amazed and delighted the teachers as they highlighted, for example, that some children who would never spontaneously greet the teachers would do so with Andy and later with the teachers in the classroom, or that some children who would never interact with their real peers in the classroom would do so with Andy.

Although successful in many respects, the implementation and the evaluation of the ECHOES environment has highlighted a number of critical issues. In what follows, we briefly discuss the main challenges we encountered in developing ECHOES as we believe that they pertain not to ECHOES specifically but, more generally, to the design and evaluation of serious games for autism. For an in-depth discussion on the technical challenges that we met and the lessons that we learned during the implementation of ECHOES, the interested reader can also consult [11].

A well-known issue, both from the point of view of the design and the evaluation of the system, concerns dealing with heterogeneous populations such as individuals with ASCs. In ECHOES, this difficulty was exacerbated by two main factors: (i) the system focuses on exploration and free play rather than strict goal-directed training; and (ii) the system was evaluated with children of different ages (range four to fourteen years) attending different schools contexts. All these factors together made it difficult to create an environment suitable for all users and greatly restricted the scope for consistent cross-participant comparisons.

Furthermore, the heterogeneity of the target population impacts on how the user model of a serious game for autism needs to be constructed. On the one hand, user modelling relies on identifying common features and behaviours across all the users for constructing the base models. On the other hand, this may not be fully feasible or even desirable for contexts where every user presents different challenges and where the domain knowledge does not lend itself to being represented in terms of *right* or *wrong* solutions. It seems that more intimate profiles for the children may need to be constructed in this context by supplementing automatic user modelling techniques with a direct involvement of teachers, parents and children themselves. To enable such a cooperative construction of the user model, the software needs to be scalable, portable and suitable for inter-context use. In the SHARE-IT project [45], an ongoing ECHOES follow-up project, we are currently exploring how such construction of user models and learning experiences for children can be achieved. In particular, an *open user model* is used to allow teachers and parents to inspect, complete and/or correct the inferences made by the model automatically [61]. While the open user model is a desirable feature for enhancing the performance of the user model, it is also important for bridging the gap between the parents' and the teachers' knowledge of the children, allowing for a more complete and shared understanding of the children abilities and needs.

Another major challenge we faced in ECHOES revolves around how to achieve a degree of sophistication for the agent personality that affords a believable and natural interaction with the children. In ECHOES, in order to cope with the complexity of modelling Andy through FATiMA, we simplified some aspects of Andy's personality with respect to its original design. However, during the evaluation of the ECHOES system, teachers and carers identified the enhancement of Andy's emotional displays as a crucial way to make the agent more credible and enticing to children as well as to introduce different levels of difficulty for children, whereby they are exposed to more subtle and varied behaviours on the part of the social partner while they progress through the sessions. Taking heed of suggestions from teachers, in SHARE-IT we have conferred a more nuanced personality on Andy by further experimenting with its affective system's parameters and developing more modulated facial expressions, body language, and voice recordings. We exploit FATiMA's full range of emotional thresholds and decay rates to allow the agent to display both positive and negative emotions as well as to change its mood within a

¹ For the full interview, see ESRC-commissioned impact video about the ECHOES project accessible on the ESRC website at <http://www.esrc.ac.uk/my-esrc/grants/RES-139-25-0395-A/outputs/read/73f03b76-a593-4d99-863e-e65fc0ab53fc>.

particular activity. In order to facilitate the modelling task, we have developed a visual *authoring tool* for manipulating and displaying the agent's emotions and automatically linking such emotional displays to the specification of the corresponding parameters in the FATiMA model (see [61] for additional details on the authoring tool).

Finally, another issue of importance to the design of serious games relates to how such technologies are used in the real-world. Thanks to our continuous interaction with teachers, practitioners and children during the course of the ECHOES project, from the design to the evaluation of the technology, we came to realise that the *intended use* of a piece of technology such as a serious game will not necessarily be reflected in its *actual use*. Therefore, the design, both at the hardware and the software level, needs to be flexible enough to allow for the actual use to be implemented successfully as needed. Our hypothesis is that such flexibility can only be achieved if the technology is both designed and evaluated in real-world contexts, as is the case with ECHOES, and if the technology is offered for use to teachers, practitioners and children, independently of the designers' and researchers' goals and regardless of their presence at the time of technology use. We are currently testing this hypothesis as part of the Shape project [44], which involves the ECHOES system and three other technologies for autism (one of which – [39] – is another serious game). These systems have been installed in four different specialist schools in the UK and their actual use is being evaluated and compared with their intended use based on the current practices employed in those schools.

Acknowledgments

The ECHOES project has been funded jointly by the Engineering and Physical Sciences Research Council (EPSRC) and Economic and Social Sciences Researcher Council (ESRC) under the TLRP-TEL programme, Grant No.: RES-139-25-0395-A. We thank all the other members of the ECHOES project (A. Alcorn, K. Avramides, J. Chen, M.E. Foster, C. Frauenberger, J. Good, K. Guldberg, W. Keay-Bright, C. Kossyvaki, O. Lemon, L. Mademtzi, R. Menzies, H. Pain, T. Rajendran, A. Waller, and S. Wass) for their contribution to the construction of the system and its evaluation. We are also very grateful to all the teachers, children and parents who have participated in the project for their insightful suggestions and their involvement in the game's evaluation.

References

- [1] B. Abirached, Y. Zhang, J. Aggarwal, B. Tamersoy, T. Fernandes, J. Miranda, V. Orvalho, 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), 2011, pp. 1–4.
- [2] A. Alcorn, H. Pain, G. Rajendran, T. Smith, O. Lemon, K. Porayska-Pomsta, M.E. Foster, K. Avramides, C. Frauenberger, S. Bernardini, Social communication between virtual characters and children with autism, in: Proceedings of the 15th International Conference on Artificial Intelligence in Education, AIED'11, Springer-Verlag, Berlin/Heidelberg, 2011, pp. 7–14.
- [3] American Psychiatric Association, Diagnostic and Statistical Manual of Mental Disorders, fourth ed., Text Revision (DSM-IV-TR), 2000.
- [4] A. Anwar, M. Rahman, S. Ferdous, S. Anik, S. Ahmed, A computer game based approach for increasing fluency in the speech of the autistic children, in: 11th IEEE International Conference on Advanced Learning Technologies (ICALT), 2011, pp. 17–18.
- [5] K. Avramides, S. Bernardini, M.E. Foster, C. Frauenberger, L. Kossyvaki, M. Mademtzi, State-of-the-art in tel to support social communication skill development in children with autism: a multi-disciplinary review, International Journal of Technology Enhanced Learning (IJTEL) 4 (5/6) (2012) 359–372.
- [6] W. Bainbridge, J. Hart, E. Kim, B. Scassellati, The benefits of interactions with physically present robots over video-displayed agents, International Journal of Social Robotics 3 (1) (2010) 41–52.
- [7] E. Barakova, G. van Wanrooij, R. van Limpt, M. Menting, 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 '07, ACM, New York, NY, USA, 2007, pp. 73–76.
- [8] S. Baron-Cohen, Mindblindness, MIT Press, Cambridge MA, 1995.
- [9] A. Battocchi, F. Pianesi, D. Tomasini, M. Zancanaro, G. Esposito, P. Venuti, A. Ben Sasson, E. Gal, P.L. Weiss, 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 '09, ACM, New York, NY, USA, 2009, pp. 197–204.
- [10] R. Beaumont, K. Sofronoff, A multi-component social skills intervention for children with asperger syndrome: the junior detective training program, Journal of Child Psychology and Psychiatry 49 (2008) 743–753.
- [11] S. Bernardini, K. Porayska-Pomsta, Planning-based social partners for children with autism, in: Proc. of the Twenty Third International Conference on Automated Planning and Scheduling (ICAPS-13), 2013.
- [12] S. Bernardini, K. Porayska-Pomsta, H. Sampath, Designing an intelligent partner for social communication in autism, in: Proc. of the Ninth Annual AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-13), 2013.
- [13] S. Bernardini, K. Porayska-Pomsta, J.T. Smith, K. Avramides, Building autonomous social partners for autistic children, in: Proc. of the 12th International Conference on Intelligent Virtual Agents (IVA-12), Lecture Notes in Computer Science, vol. 7502, 2012, pp. 46–52.
- [14] A. Bosseler, D. Massaro, 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) (2003) 653–672.
- [15] S. Cobb, Virtual environments supporting learning and communication in special needs education, Topics in Language Disorders 27 (3) (2007) 211–225.
- [16] K. Dautenhahn, I. Werry, Towards interactive robots in autism therapy: background, motivation and challenges, Pragmatics and Cognition 12 (1) (2004) 1–35.
- [17] M. Davis, N. Otero, K. Dautenhahn, C. Nehaniv, S. Powell, 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, 2007 (ICDL 2007), July 2007, pp. 64–69.
- [18] G. Dawson, S. Rogers, J. Munson, M. Smith, J. Winter, J. Greenson, A. Donaldson, J. Varley, Randomized, controlled trial of an intervention for toddlers with autism: the early start denver model, Pediatrics 125 (1) (2010) 17–23.
- [19] S. De Freitas, Learning in immersive worlds, Tech. rep., Joint Information Systems Committee, Bristol, 2006. <<http://www.jisc.ac.uk/elioutcomes.html>> (accessed 20.10.12).
- [20] J. Dias, A. Paiva, Feeling and reasoning: a computational model for emotional characters, in: Progress in Artificial Intelligence, Lecture Notes in Computer Science, vol. 3808, Springer, Berlin/Heidelberg, 2005, pp. 127–140.

- [21] S.L. Finkelstein, A. Nickel, L. Harrison, E.A. Suma, T. Barnes, 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*, 2009, pp. 249–250.
- [22] C. Frauenberger, J. Good, W. Keay-Bright, Phenomenology, a framework for participatory design, in: *Proceedings of the 11th Biennial Participatory Design Conference*, PDC '10, ACM, New York, NY, USA, 2010, pp. 187–190.
- [23] M. Frutos, I. Bustos, B. Zapirain, A. Zorrilla, Computer game to learn and enhance speech problems for children with autism, in: *16th International Conference on Computer Games (CGAMES)*, 2011, pp. 209–216.
- [24] O. Golan, E. Ashwin, Y. Granader, S. McClintock, K. Day, V. Leggett, S. Baron-Cohen, A multi-component social skills intervention for children with asperger syndrome: the junior detective training program, *Journal of Autism and Developmental Disorders* 40 (2010) 269–279.
- [25] M. Goodwin, Enhancing and accelerating the pace of autism research and treatment: the promise of developing innovative technology, *Focus on Autism and Other Developmental Disabilities* 23 (2) (2008) 125–128.
- [26] O. Grynspan, J.-C. Martin, J. Nadel, Multimedia interfaces for users with high functioning autism: an empirical investigation, *International Journal on Human-Computer Studies* 66 (8) (2008) 628–639.
- [27] H. Hoffman, Virtual reality therapy, *Scientific American Magazine* (2004) 58–65.
- [28] I.M. Hopkins, M.W. Gower, T.A. Perez, D.S. Smith, F.R. Amthor, 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) (2011) 1543–1555.
- [29] M.E. Hoque, J.K. Lane, R.E. Kaliouby, M.S. Goodwin, R.W. Picard, Exploring speech therapy games with children on the autism spectrum, in: *Proceedings of the 10th Annual Conference of the International Speech Communication Association, INTERSPEECH 2009*, 2009, pp. 1455–1458.
- [30] G. Iarocci, J. McDonald, Sensory integration and the perceptual experience of persons with autism, *Journal of Psychiatric Research* 3 (2006) 181–197.
- [31] C. Kidd, C. Breazeal, Effect of a robot on user perceptions, in: *Proceedings. 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 4, 2004, pp. 3559–3564.
- [32] H. Kozima, M. Michalowski, C. Nakagawa, Keepon: a playful robot for research, therapy, and entertainment, *International Journal of Social Robotics* 1 (1) (2009) 3–18.
- [33] M. Kriegel, R. Aylett, Emergent narrative as a novel framework for massively collaborative authoring, in: *Proceedings of the 8th International Conference on Intelligent Virtual Agents*, 2008, pp. 73–80.
- [34] M.Y. Lim, J. Dias, R. Aylett, A. Paiva, Creating adaptive affective autonomous NPCs, *Journal of Autonomous Agents and Multi-Agent Systems* 24 (2) (2012) 287–311.
- [35] O.I. Lovaas, Behavioral treatment and normal educational and intellectual functioning in young autistic children, *Journal of Consulting and Clinical Psychology* 55 (1) (1987) 3–9.
- [36] D. Massaro, Embodied agents in language learning for children with language challenges, in: *Proceedings of the 10th International Conference on Computers Helping People with Special Needs, ICCHP'06*, Springer-Verlag, Berlin/Heidelberg, 2006, pp. 809–816.
- [37] ASC-Inclusion, 2012. <<http://asc-inclusion.eu/>> (accessed 20.10.12).
- [38] Brigadoon, John Lester, 2012. <<http://braintalkblogs.com/brigadoon>> (accessed 20.10.12).
- [39] Cospatial, 2012. <<http://cospatial.fbk.eu/>> (accessed 20.10.12).
- [40] ECHOES, 2013. <<http://echoes2.org/>> (accessed 22.07.13).
- [41] Medical Research Council, Autism – Research Review, 2001. <<http://www.mrc.ac.uk/Utilities/Documentrecord/index.htm?d=MRC002394>> (accessed 20.10.12).
- [42] National Autistic Society, What is autism? 2012. <<http://www.autism.org.uk/about-autism.aspx>> (accessed 20.10.12).
- [43] Second Life, 2012. <<http://secondlife.com/>> (accessed 20.10.12).
- [44] Shape, 2012. <<http://www.birmingham.ac.uk/research/activity/education/shape>> (accessed 15.07.13).
- [45] SHARE-IT, 2013. <<http://shareitproject.wordpress.com/>> (accessed 15.07.13).
- [46] TurboSquid, Andy, 2012. <<http://www.turbosquid.com/3d-models/human-kid-child-3d-model/544104>> (accessed 20.10.12).
- [47] M. Milne, M. Luerssen, T. Lewis, R. Leibbrandt, D. Powers, Development of a virtual agent based social tutor for children with autism spectrum disorders, in: *Proc. of the International Joint Conference on Neural Networks*, 2010, pp. 1–9.
- [48] P. Mitchell, S. Parsons, A. Leonard, Using virtual environments for teaching social understanding to 6 adolescents with autistic spectrum disorders, *Journal of Autism and Developmental Disorders* 37 (3) (2007) 589–600.
- [49] M. Moore, S. Calvert, Brief report: vocabulary acquisition for children with autism: teacher or computer instruction, *Journal of Autism and Developmental Disorders* 30 (4) (2000) 359–362.
- [50] D. Murray, Autism and information technology: therapy with computers, in: D.F. Publishers (Ed.), *Autism and Learning: A Guide to Good Practice*, Brookes, 1997, pp. 100–117.
- [51] T. Myers, L. Swan-Kremeier, S. Wonderlich, K. Lancaster, J. Mitchell, The use of alternative delivery systems and new technologies in the treatment of patients with eating disorders, *International Journal of Eating Disorders* 36 (2) (2004) 123–143.
- [52] H.A.M. Noor, F. Shahbodini, N.C. Pee, Serious game for autism children: review of literature, in: *Proceedings of the International Conference on Computer Games, Multimedia and Allied Technology (ICCGMAT '12)*, 2012.
- [53] A. Ortony, G.L. Clore, A. Collins, *The Cognitive Structure of Emotions*, Cambridge University Press, 1988.
- [54] M. Pantic, R. Cowie, F. D'Ercole, D. Heylen, M. Mehu, C. Pelachaud, I. Poggi, M. Schroder, A. Vinciarelli, *Social Signal Processing: The Research Agenda* (2011) 511–538.
- [55] M. Pantic, A. Nijholt, A. Pentland, T. Huang, Human-centred intelligent human computer interaction (HCI2): how far are we from attaining it?, *International Journal on Autonomous and Adaptive Communications Systems* 1 (2) (2008) 168–187.
- [56] M. Pantic, L.J.M. Rothkrantz, Towards an affect-sensitive multimodal human–computer interaction, *Special Issue on Multimodal Human–Computer Interaction (HCI)* 91 (2003) 1370–1390.
- [57] S. Parsons, S. Cobb, State-of-the-art of virtual reality technologies for children on the autism spectrum, *European Journal of Special Needs Education* 26 (3) (2011) 355–366.
- [58] S. Parsons, P. Mitchell, The potential of virtual reality in social skills training for people with autistic spectrum disorders, *Journal of Intellectual Disability Research* 46 (5) (2002) 430–443.
- [59] S. Parsons, P. Mitchell, A. Leonard, The use and understanding of virtual environments by adolescents with autistic spectrum disorders, *Journal of Autism and Developmental Disorders* 34 (4) (2004) 449–466.
- [60] S. Parsons, P. Mitchell, A. Leonard, Do adolescents with autistic spectrum disorders adhere to social conventions in virtual environments?, *Autism* 9 (1) (2005) 95–117.
- [61] K. Porayska-Pomsta, K. Anderson, S. Bernardini, K. Guldberg, T. Smith, L. Kossivaki, S. Hodgings, I. Lowe, Building intelligent authorable serious game for autistic children and their carers, in: *Proceedings of the 10th International Conference on Advances in Computer Entertainment (ACE 2013)*, 2013.
- [62] K. Porayska-Pomsta, C. Frauenberger, H. Pain, G. Rajendran, T. Smith, R. Menzies, M.E. Foster, A. Alcorn, S. Wass, S. Bernardini, K. Avramides, W. Keay-Bright, J. Chen, A. Waller, K. Guldberg, J. Good, O. Lemon, Developing technology for autism: an interdisciplinary approach, *Personal Ubiquitous Computing* 16 (2) (2011) 117–127.
- [63] B. Prizant, A. Wetherby, E. Rubin, A. Laurent, The scerts model: a transactional, family-centered approach to enhancing communication and socioemotional ability in children with autism spectrum disorder, *Infants and Young Children* 16 (4) (2003) 296–316.
- [64] B. Prizant, A. Wetherby, E. Rubin, A. Laurent, P. Rydell, *The SCERTS® Model: A Comprehensive Educational Approach for Children with Autism Spectrum Disorders*, Brookes, 2006.
- [65] C. Putnam, L. Chong, Software and technologies designed for people with autism: what do users want? in: *10th International ACM SIGACCESS Conference on Computers and Accessibility*, 2008, pp. 3–8.

- [66] M.M. Rahman, S. Ferdous, S.I. Ahmed, A. Anwar, Speech development of autistic children by interactive computer games, *Interactive Technology and Smart Education* 8 (4) (2011) 208–223.
- [67] M. Riedl, C.J. Saretto, R.M. Young, Managing interaction between users and agents in a multi-agent storytelling environment, in: *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '03*, ACM, New York, NY, USA, 2003, pp. 741–748.
- [68] S.J. Rogers, L.A. Vismara, Evidence-based comprehensive treatments for early autism, *Journal of Clinical Child and Adolescent Psychology* 37 (1) (2008) 8–38.
- [69] S. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, second ed., Prentice Hall, 2003.
- [70] C.A. Smith, R.S. Lazarus, Emotion and adaptation, in: I. Vlahavas, D. Vrakas (Eds.), *Handbook of Personality: Theory and Research*, Guilford, New York, 1990, pp. 609–637. Chapter 23, L.A. Pervin (Ed.).
- [71] A. Tartaro, J. Cassell, Using virtual peer technology as an intervention for children with autism, in: J. Lazar (Ed.), *Universal Usability: Designing Computer Interfaces for Diverse User Populations*, John Wiley and Sons, Ltd., New York, 2006, pp. 231–262. Chapter 8.
- [72] A. Tartaro, J. Cassell, Playing with virtual peers: bootstrapping contingent discourse in children with autism, in: *Proc. ICLS 2008*, 2008, pp. 382–389.
- [73] J. Wainer, D. Feil-Seifer, D. Shell, M. Mataric, The role of physical embodiment in human-robot interaction, in: *The 15th IEEE International Symposium on Robot and Human Interactive Communication*, 2006, pp. 117–122.
- [74] M. Walker, A. Armfield, What is the makaton vocabulary?, *Special Education: Forward Trends* 8 (3) (1981) 19–20.
- [75] L. Wiggs, G. Stores, Sleep patterns and sleep disorders in children with autistic spectrum disorders: insights using parent report and actigraphy, *Developmental Medicine and Child Neurology* 46 (2004) 372–380.
- [76] C. Williams, B. Wright, G. Callaghan, B. Coughlan, Do children with autism learn to read more readily by computer assisted instruction or traditional book methods?, *Autism* 6 (1) (2002) 71–91.
- [77] L. Wing, *The Autism Spectrum*, Constable, London, 1996.
- [78] D. Wood, J. Murphy, K. Center, C. Russ, R. McLay, D. Reeves, J. Pyne, R. Shilling, J. Hagan, B. Wiederhold, Combat related post-traumatic stress disorder: a multiple case report using virtual reality graded exposure therapy with physiological monitoring, *Studies in Health Technology and Informatics* 132 (2008) 556–561.
- [79] M. Wooldridge, N.R. Jennings, Intelligent agents: theory and practice, *Knowledge Engineering Review* 10 (2) (1995) 115–152.
- [80] G. Yannakakis, J. Hallam, Modeling and augmenting game entertainment through challenge and curiosity, *International Journal on Artificial Intelligence Tools* 16 (6) (2007) 981–999.