A Tabletop Interactive Storytelling System: Designing for Social Interaction

Thijs Alofs

Human Media Interaction
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands
E-mail: t.alofs@gmail.com

Mariët Theune

Human Media Interaction
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands
E-mail: m.theune@utwente.nl

Ivo Swartjes

Ranj Serious Games
Lloydstraat 21m
3024 EA Rotterdam
The Netherlands
E-mail: ivo@ranj.nl

Abstract: This paper presents the Interactive Storyteller, a multi-user interface for AI-based interactive storytelling, where stories emerge from the interaction of human players with intelligent characters in a simulated story world. To support face-to-face contact and social interaction, we position users around a shared multi-touch table, which very much resembles the social setting of traditional tabletop board games. To our knowledge, the Interactive Storyteller is the first AI-based interactive storytelling system that combines an emergent narrative approach with the social aspects of traditional tabletop board games. We carried out user experiments to investigate to what extent our system supports social interaction. By analysing the interactions of pairs of children with the Interactive Storyteller, we determined which system aspects triggered cooperation and highly social behaviour and which aspects caused players to behave less socially. We also tried to find out whether the use of tangible playing pieces offered any advantages over touch-only interaction, but we did not find any differences between the two.

Biographical notes: Thijs Alofs is a developer and interaction designer at Hesitate creative concepting. Prior to this he worked at Montix as a frontend developer. His interests include human-computer interaction, touch interfaces, usability, and Android. The research presented in this article was performed as the final assessment in obtaining his masters degree. Thijs holds a bachelors degree in Computer Science and a masters degree in Human Media Interaction, both from the University of Twente, The Netherlands.

Mariët Theune is assistant professor at the Human Media Interaction research group of the University of Twente (Enschede, The Netherlands). Since 2004 she has been coordinating the research on digital interactive storytelling at HMI, currently as part of the COMMIT project "Interaction for Universal Access". Her other research interests include natural language generation, multimodal information presentation, embodied conversational agents and other natural language dialogue systems. She holds a master's degree in computational linguistics from Utrecht University, and a PhD in natural language generation from the Eindhoven University of Technology, the Netherlands.

Ivo Swartjes is a senior programmer at Ranj Serious Games, where he acts as technical lead in serious games projects. Next to the design and implementation of game software architectures, he develops and improves their underlying engines and tools to optimise the game development pipeline and to enable new game formats. He also investigates the use of AI models for dialogues with non-player characters and especially their implications for the player experience, project size and production pipeline. Prior to Ranj Serious Games, he worked as a postdoctoral researcher at the University of Twente. He holds a masters degree in Computer Science and a PhD in Human Media Interaction from the University of Twente, the Netherlands.

1 Introduction

Most models of what constitutes an enjoyable, playful experience include social interaction as an important component (Choi & Kim 2004, Csikszentmihalyi 1991, Korhonen et al. 2009). Social elements such as cooperation, collaboration, and competition tend to increase the intrinsic motivation to engage in and maintain

some activity (Malone & Lepper 1987) and may even be the main motivational element of a game (Karadimitriou & Roussou 2011).

In this paper we discuss the design and evaluation of a multi-user tabletop interface for an interactive storytelling system that was designed to support social interaction. In the *Interactive Storyteller* users can control the actions of one or more characters in a simulated story world. Each user can play one character, but collaborative control is also possible, with users cooperating to make decisions for the characters. From the ongoing interaction between the characters (which can be either player- or system-controlled) and through their choices of action a story emerges; this approach to interactive storytelling is called 'emergent narrative' (Aylett et al. 2005).

By using a multi-touch table we aim to achieve interactive storytelling that resembles the social setting of a tabletop board game, stimulating face-to-face contact and social interaction. Like existing digital tabletop board games (Magerkurth et al. 2004, Mandryk & Maranan 2002, Piper et al. 2006), we try to combine the dynamics and intelligence of computer games with the social advantages of traditional board games. To reinforce the resemblance with board games, we investigate the use of tangible playing pieces that represent characters for physical interaction.

The idea of using tabletop interfaces for digital storytelling is not entirely new. Several tabletop interfaces exist for storytelling systems (Alves et al. 2010, Cappelletti et al. 2004, Helmes et al. 2009), but they only focus on facilitating storytelling, trying to stimulate collaboration and creativity. Unlike our system, they do not use artificial intelligence (AI) to contribute to the story.

The interfaces of current AI-based interactive storytelling systems are mostly like those of computer games, where a single user interacts from a first person perspective with 2D or 3D virtual characters on a computer screen (Aylett et al. 2005, Mateas & Stern 2003, Pizzi & Cavazza 2007). A few systems focus on collaborative storytelling by multiple users, but these systems also have interfaces similar to computer games (Kriegel et al. 2010, Prada et al. 2002). To our knowledge, our Interactive Storyteller is the first AI-based storytelling system with a tabletop interface.

The technical framework underlying the Interactive Storyteller is a multi-agent system for story generation, in which intelligent agents act out the role of characters in the story. These agents can plan and execute sequences of actions to satisfy their character's goals, taking into account the current state of the story world, and the mental state of the character. Actions are performed in turns, similar to turn-taking in many traditional board games. For more information on the storytelling framework, see Swartjes & Theune (2008), Swartjes (2010).

Two story domains are currently available in our storytelling framework: one about pirates, and one based on the "Little Red Riding Hood" (LRRH) story (Swartjes & Theune 2009). We used the latter domain for our experiments with the Interactive Storyteller, because we considered it to be more coherent, easier to understand for new users, and easier to visualise than a pirate ship. Given that LRRH is a children's fairytale and because children are an interesting target group for interactive storytelling, we decided to focus our research on children aged 8 to 11. However, a goal kept in mind was that the interface should also be appropriate for adults (when used with more adult domains).

Below, we first give a brief overview of related work in Section 2. Then, in Section 3 we describe the interface design of the Interactive Storyteller and its underlying considerations. In Section 4 we present our observations from a small-scale user evaluation, focusing on the question which factors stimulated or hindered social interaction. We end with a discussion and some pointers to future work.

2   Related Work

In this section we briefly discuss some work on digital tabletop board games that formed the main inspiration for the interface design of the Interactive Storyteller.

STARS is a platform for developing computer augmented tabletop games (Magerkurth et al. 2003, 2004). It consists of a touch table on which physical playing pieces can be placed. These objects are recognised by an overhead camera. An RFID antenna below the table's surface provides the STARS platform with another way to receive information from the physical world (Magerkurth et al. 2003). Besides the horizontal touch screen, a large vertical display is used for additional images or video. Every user of the STARS platform can use a personal digital assistant (PDA) for personal information and additional interaction using voice input and output. Earphones connected to the PDA can provide private audio information. There are also speakers for public sounds, for instance ambient audio samples or atmospheric music. Several games have been developed for the STARS platform so far, including the role playing game KnightMage and an adaptation of Monopoly (Magerkurth et al. 2004).

The research by Tse et al. (2007) demonstrates the application of a multi-user, multi-touch interface on existing single player computer games, using the MERL DiamondTouch hardware (Dietz & Leigh 2001). Speech input and gesture input from the multi-touch table have been mapped to the games Warcraft III and The Sims with the Gesture Speech Infrastructure (Tse et al. 2006). As in our Interactive Storyteller, the view and interaction in both games is mainly map-based. The research by Tse et al. shows that a tabletop interface can provide a social setting for game interaction. There are no individual views and interfaces, everybody shares a common view and has the interface within reach. This provides a social setting for collaborative play.

The SurfaceScapes team (Master students from the Entertainment Technology Center at Carnegie Mellon University) developed a prototype for a Dungeons & Dragons (DnD) experience on the Microsoft Surface touch table (Coldewey 2010). An interesting video demonstration of this prototype can be watched at `http://vimeo.com/7132858`. There are two kinds of maps in the SurfaceScapes game. The world map is movable, zoomable and rotatable. After selecting a location on the world map, a fixed battle map with a grid is shown. On this grid the physical character objects are placed and virtual objects are displayed.

Other digital tabletop board games related to the Interactive Storyteller are False Prophets by Mandryk & Maranan (2002), who were among the first to explore hybrid video/board games, and SIDES by Piper et al. (2006).

3 Interface Design

The interface design of the Interactive Storyteller is meant to be generic and easily usable for any existing or future story domain. The only essential additions to enable an existing specification of a story domain to be used in the Interactive Storyteller are images and coordinates. For any domain, pictures that represent the characters and a map of the story world have to be added. The coordinates of the locations on the map, which link the story world specification to its visual representation, have to be provided. Other necessary domain-specific information for the interface can be retrieved from the specification of the story world.

We use a multi-touch table based on infrared reflection that is capable of identifying tangible objects through fiducial markers. The MT4j framework (Multi-touch for Java, `http://www.mt4j.org/`) is used for multi-touch support.

Figure 1   The system at the time of the early user tests (see Section 3.4).

3.1 Story World and View

The LRRH story domain is a variation on the original story of Little Red Riding Hood. It contains three characters (Red, Grandma, and Wolf) and five locations (Red's house, Grandma's house, the clearing in the forest, the lake, and the beach). Typical actions for a character currently available in the LRRH domain

are amongst others: walking, greeting someone, stealing things, crying, eating something, baking a cake, and poisoning food. Characters can plan a series of such actions to try to achieve a goal they have. For instance, for the goal of 'Red' wanting to poison the wolf, a possible plan might involve baking a cake, poisoning it with cyanide, and giving it to Wolf (expecting him to eat it). To increase the number of character actions that could be selected by the players, we added a few actions to the LRRH domain that are local in nature and have no effects on the story world. These include actions such as watching birds, enjoying the sun, taking a nap, etc. In Section 4.5 we show an example story in the LRRH domain that was created during the user experiments we carried out.

The visualisation of the story world is presented to users and possible spectators on a shared visual surface. Just as with traditional board games, it is important that people on all sides of the table have a similar view on the story world, therefore we chose a top-down map view. The story locations are marked by blue circles on the map. For aesthetic reasons, we decided to draw the characters and houses on the LRRH map not strictly from their top-side view, but from a more recognisable angle. However, to prevent one side of the table from being optimal for perceiving the story world, characters are displayed in one direction and houses are projected the other way. Another solution for this orientation issue is autorotation, as was used in KnightMage (Magerkurth et al. 2004). Autorotation is however only useful when the physical position of the user who currently has to choose an action can be determined, which is not the case with our hardware.

3.2 Physical Interaction

Like in the systems KnightMage (Magerkurth et al. 2004) and False Prophets (Mandryk & Maranan 2002), in one version of our system users can change the locations of characters by moving physical toys that represent the story characters across the surface of the multi-touch table. These tangibles provide tactile interaction that is expected to be intuitive because it strongly resembles the interaction offered by many familiar board games. Moreover, "tangible objects can invite us to interact by appealing to our sense of touch, providing sensory pleasure and playfulness" (Hornecker & Buur (2006), p. 440). To see whether tangibles really provide added value for the Interactive Storyteller, we also developed a touch-only version of the interface where the characters have graphical representations that can be moved around by dragging them across the table surface.

In our storytelling system locations are always discrete: characters are at one location, or the next, but never half-way in between. When a user moves a tangible to an adjacent location, the blue circle of the destination location turns green to indicate that this action is allowed. When the user moves the character to a location that is not in direct reach of the character's current location, the circle of the destination turns red to indicate that this action is not allowed.

Users might put tangibles outside the circles that mark locations on the map. The system is unable to physically move tangibles away from such non-locations. Interventions, like the system asking the user to move a tangible to a particular location, are not used because they distract the user from the story. This means the system has to be able to deal with tangibles being anywhere on the map. When

a tangible is not at its character's actual location in the story world, a (dotted) blue line is shown connecting the tangible to its character's location. In the touch-only version, the system is in principle capable of automatically moving character images back to their correct location. However, we decided not to implement this feature; we wanted to keep the touch-only version as similar as possible to the touch+tangible version of the system to ensure a 'fair' comparison in our experiments.

3.3 Interface Elements

For the selection of non-move actions by users, there is an *Action Selection Interface* (ASI). An important requirement we had for this interface was that it should be quickly usable for any storytelling domain. A very specific and intuitive (iconographic) ASI can be developed by focussing on one particular domain to fit the users' needs in that particular virtual world. This approach is commonly used in computer games, but we considered it more important for the interface to stay generic, and not designed for a specific story domain and set of actions. Therefore, we decided to refrain from icon-based or other graphics-based ways for action selection, and to use a flexible text-based approach instead.

Every time a character gets the turn, the knowledge base retrieves the set of all actions that are possible in the story world for that character, at that specific time and location. The ASI can be seen in the centre of Figure 1 and in more detail in Figure 2. Because of the young target user group, all text is displayed in our native language Dutch. The user first selects a category in the centre bar of the ASI and then an action within that category. After that, the round confirmation button in the centre bar is enabled and can be used to confirm the selected action and pass the turn to the next character.
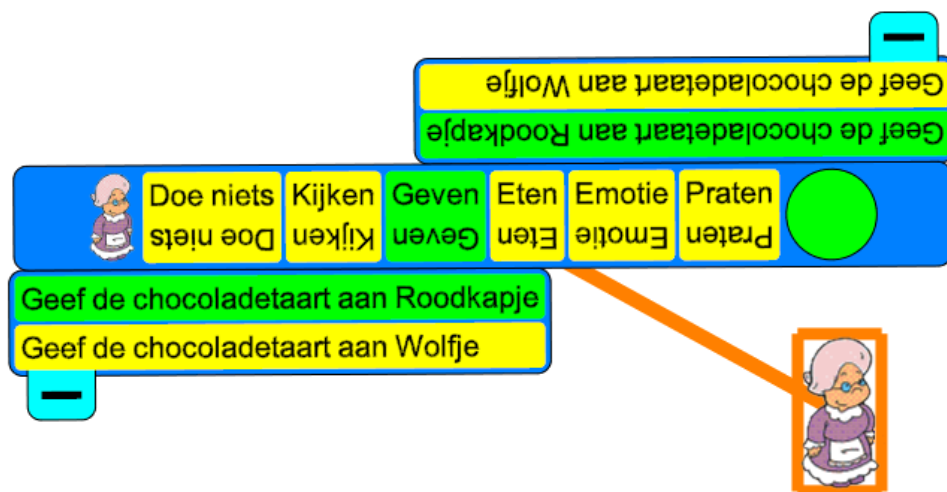

Figure 2 The Action Selection Interface, connected to character that has the turn.

Users and spectators can read the results of actions that characters perform in the *story areas*: the two scrolls that can be seen on the tabletop display in Figure 1 and in the screenshot of the final interface design in Figure 4. If users want more or fewer lines of text to be visible in a story area, this goal can be achieved by touching the top end of the scroll and rolling it up or down.

The ASI and the story areas occlude the view of the part of the map behind them. A balance has to be found between good visibility of the map and its contents, and the readability of the ASI and story areas. To achieve this subjective balance, we decided to keep the user in control of the size and placement of the ASI and the story areas. The user can move these elements around by dragging them with a finger. By dragging with two fingers at the same time, it is possible to rotate and resize them. The user can choose to find a static arrangement that generally works well in a particular story world, or keep changing sizes and arrangements depending on the current state of the story and places of interest. This approach is in line with the recommendations of Kruger et al. (2003) that users should always have the option to rotate items freely.

Because we consider it to be important that users or spectators from all sides of the table have an equal view, all text in the ASI is presented in two directions instead of one. When users or spectators are standing on all four sides of the multi-touch table, the optimal layout is to position the ASI under an angle of 45 degrees with the sides of the table. We expect this angle to be acceptable for most readers. Having a shared ASI saves much space compared to having separate control areas on all four sides of the tabletop for different users, as in the SIDES system (Piper et al. 2006). Moreover, if everybody is standing on one side of the tabletop, the text at the opposite side of the ASI can be hidden by touching the minus symbol on that side, conserving even more screen estate.

3.4 Early User Tests

During the interface design phase we performed some informal user tests. These tests involved five test subjects (three boys aged 8, 8 and 10, and two girls aged 10 and 11) who interacted with an early prototype of the system. We found that despite the limited graphics, the children were engaged by the system and enjoyed playing with it. With only a very limited explanation they understood how to interact with the system. Several children discussed possible actions together and some even planned a sequence of actions to pursue a particular storyline.

Nevertheless, we did observe some small problems that led us to make several improvements to the system. For example, we found that the children's fingertips were often badly recognised because they were very small, while at the same time the rest of the hand did get recognised while hovering above the surface (see Figure 3). By fine-tuning some recognition parameters we managed to reduce these issues.

Another problem we observed was that although the children had been told to use the story areas to keep track of what was happening in the story (in particular, the actions performed by other characters), some of the children never looked at the story areas. This behaviour often resulted in these users ending up confused and less immersed in the story. To address this issue, we decided to offer the same information in another (complementary) modality by vocalising it with Loquendo
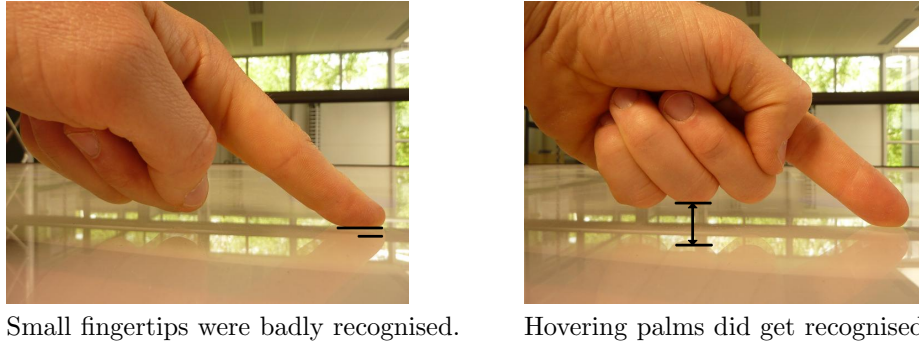
Small fingertips were badly recognised.   Hovering palms did get recognised.

Figure 3   Recognition problems in the setup used during the early user tests.

text-to-speech, while keeping the story areas as a time-independent source of the same information about the story. After introducing audio as a new modality, we decided to also allow the addition of action specific sounds. Although domain specific, associating actions with sounds is a quick and easy way to present audible feedback of an action while at the same time enriching the user experience by stimulating what has been termed 'sensory curiosity' (Malone & Lepper 1987).

Finally, we discovered that users often lost track of turn-taking. We solved this issue by connecting the ASI by a solid line to the character that has the turn, to make clear to which character the actions in the ASI belong (see Figure 2).

3.5   Pilot User Test

A pilot user test involving two adult users was performed to ensure that the system was ready for our evaluation experiment with children, described in Section 4. The pilot user test (as well as the final evaluation experiment) was done using more advanced hardware than used in the early user tests. This better hardware resulted in a severe reduction of the number of technical problems. The pilot user test did reveal a few remaining problems with the interface. These were small and could easily be fixed.

First, it turned out that the turn-taking solution described in the previous section was still not sufficient. Therefore we added a small icon depicting the active character to the ASI to make the link with the character even more explicit (see Figure 2). Another change we made based on the pilot user test was to establish minimum and maximum sizes for every interface element. This was done to keep the users from making the elements so small or big that they could no longer be effectively manipulated. We also increased the size of the blue location circles on the map to provide more space for multiple images or tangibles.

Finally, we changed the ASI so that it displayed a message whenever the system-controlled character, Wolf, was busy making a plan for longer than one second. System-controlled characters use partial-order planning (Russell & Norvig 1995) to reach their goals. Exponential computing time is a known issue of partial-order planning algorithms. For example in the LRRH domain, Wolf is usually hungry, and then assumes the goal to eat something. When there are no cakes (the only edible objects) anywhere in the story world, the planner continues until
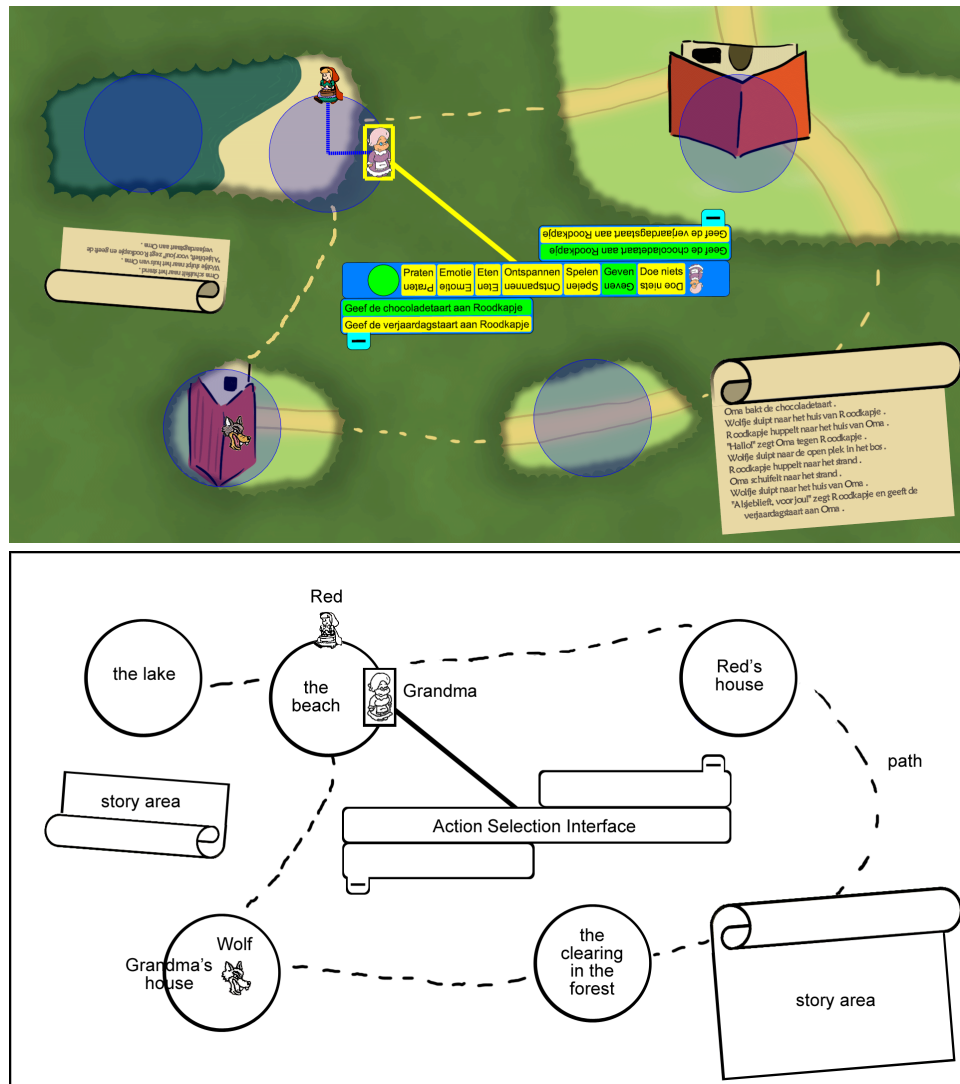
Figure 4   Screenshot of the final interface design. The schematic version shows the functional roles of the various interface elements.

the maximum search depth is reached, because there is no possible plan to satisfy Wolf's hunger. Reaching this maximum search depth takes 18 seconds on average, but occasionally even double that time. Lowering the maximum search depth to decrease the planning time was not an option, because of the minimal number of steps required to achieve the "eat-something" goal.

Figure 4 shows the final version of the interface after making the last changes based on the pilot test results.

4 Evaluation

The Interactive Storyteller was designed to support a social form of interactive storytelling, facilitating group play as opposed to solitary play. To investigate whether we achieved this goal, we carried out a small-scale, qualitative user evaluation.

We adopted the board game metaphor to establish a social setting for multi-user interactive storytelling. To reinforce the resemblance with traditional board games, we enabled the use of tangible playing pieces that represent the characters in a story. As discussed earlier, the physical properties of tangible playing pieces limit the flexibility of the system, in contrast to an image-based representation of characters. To find out whether tangibles offer advantages that might compensate for this reduced flexibility, we created a touch-only version of the Interactive Storyteller that uses images instead of tangibles to represent characters. We let the participants interact with both versions of the interface in order to investigate the differences in their experience and observed social behaviour.

4.1 Procedure

We invited four pairs of 8-11 year old children to play with the system. None of the children had interacted with the Interactive Storyteller before. All children were pairs of siblings or friends. Every pair interacted with both versions of the system, in a counterbalanced order. Table 1 provides an overview of the participants.

	Participants (age)	Relationship	Session 1	Session 2
Pair 1	Female (9) + male (9)	friends	touch-only	tangibles
Pair 2	Male (8) + female (10)	siblings	tangibles	touch-only
Pair 3	Male (8) + male (10)	siblings	touch-only	tangibles
Pair 4	Female (10) + female (11)	friends	tangibles	touch-only

Table 1 The four pairs of participants in our user evaluation.

The procedure was as follows. First, the facilitator (the researcher conducting the evaluation) gave the children a brief introduction to the experiment, explaining to them that they were helping us to test a new system, and that their feedback would be used to improve it. Our specific interest in social interaction was not mentioned. Then, the children received a brief instruction on how to use the interface. They were told that they could use the system to create a story, but were not given a specific goal to achieve, because we did not want to steer them in any specific direction. They were not asked to talk or think aloud, because we were interested in their spontaneous behaviour while interacting with the system.

Due to the generic design of the Interactive Storyteller, every character can be set up as being controlled by a user or by an AI-based agent. We decided that in our experiments always one character should be controlled by AI and two by the users. This allowed each child to have their own character, increasing the resemblance with board games where each player usually has their own playing piece. The character controlled by AI should always be the same in order to avoid

any character-specific variations in the results. Being the antagonist and the only non-human character, we decided that Wolf would not be controllable by the human test subjects and would always be controlled by AI.

The child located nearest to Red's house got assigned the character of Red to control; the other child controlled the actions of Grandma. If one or both of the children had a preference for a character, they were instructed to position themselves on the corresponding side of the table.

After the introduction and instructions, the children interacted with the first version of the system. Afterwards, the children were asked a few short questions: they were asked what they thought was good about the system, what was bad, and what they thought about the way they could move their characters. These questions were followed by an interaction session with the other version of the system, after which they were asked the same three questions about this second version.

The final part of the evaluation took the form of an informal interview with some more questions. Four questions focused on social interaction, asking the children in which version they thought they had most often cooperated, talked, laughed and had eye contact. They were also asked which version was more fun, which version was easiest to operate, and if they would like to play again with one of the two versions. Because we were interested in their attitudes toward group play, we asked the children if they enjoyed playing with two people, if they would also want to play the game alone, if they would like to play the game with more players and how many. Finally, we asked if they had any experience with touch devices such as modern smartphones or tablets such as the iPad, which might have influenced their ease of interaction with, and expectations of, the multi-touch interface.

Both the facilitator and the accompanying parent were in the room with the children while they were playing with the system. The former was present to provide technical assistance when necessary, and the latter to make both parent and children feel at ease during the experiment. However, the parent's presence did occasionally lead to the children seeking attention from the parent (see Section 4.3.2).

The interaction sessions were ended by the facilitator after or just before 15 minutes of video recording. Although the children continued without encouragement until they were stopped, in most cases the children did not spend the full 15 minutes interacting with the system. Some interaction time was lost due to questions, technical problems, talk to one of the adults, etc. After subtracting these non-interaction times, we found an average actual interaction time of 12 minutes (min. 9, max. 14 minutes). Each interaction session consisted of one story except when a crash required a restart, which happened twice.

The entire experiment, including the interview sessions, was recorded by video cameras on both short sides of the multi-touch table. The users stood on the opposing long sides of the table; see Figure 5. The recordings enabled easy observation of the children's interaction with the system, the state of the story world, facial expressions, social behaviour, etc.

Figure 5   Experiment setup and interaction.

4.2   Annotation

To see to what extent the interface of the Interactive Storyteller stimulated social interaction, we annotated the video recordings of the interaction sessions. For annotation we used the Social Play Continuum (SPC) observation scheme, because of its explicit focus on social play and social interaction (Broadhead 2004). The SPC offers a small number of interaction categories that cover the entire spectrum from non-social to highly social behaviour (Broadhead 2006, 2009). This is in contrast with other play annotation schemes such as the Play Observation Scale (POS), which tend to focus more on annotation of single behaviours (Rubin 1982). The POS is also less suitable because it pays no special attention to social behaviour, which is the main focus of our research.

The SPC distinguishes four classes of social play (or 'play domains') with increasing levels of reciprocity and momentum in language as well as actions: associative play, social play, highly social play and cooperative play. In the original SPC, parallel play (children playing at the same time and in the same location, but each by themselves) is part of the least social domain: associative play. However, we considered parallel play as less social than other associative play behaviours such as looking at peers, self-talk, and imitation. Therefore, we added a category named "Different focus of attention" on the non-social side of the spectrum. This category includes, but is not limited to parallel play. This new category allowed us to mark those parts of the interaction where the users paid no attention at all to each other. Including this additional category, the annotated play domains and their key behaviours we used in our research are:

1. **Different focus of attention**: parallel play, moments where players are paying no attention to each other at all (for example when distracted).

2. **Associative play**: watching others play, self talk, looking at peer, imitation, no interaction.

3. **Social play**: smiling, laughter, play noises, play voices, mutual eye contact, isolated verbal interactions.

4. **Highly social play**: dialogue, reciprocal sequences, clustered social play (eye contact with laughter), development of joint play themes.

5. **Cooperative play**: offering and accepting help, problem solving, role play, achieving shared goals.

As a basis for the annotation, we used the SPC tables by Pat Broadhead (Broadhead 2009). These tables describe the SPC domains on different levels, global and detailed. The table describing detailed behaviours proved very useful during annotation as it could be used to map observed actions and language to the corresponding SPC domains. Annotation was performed using the Anvil video annotation tool (Kipp 2001).

The annotators put every video frame of actual interaction time (where the system was running smoothly and the facilitator was not talking) in a segment, and classified it in one of the 5 categories of the extended SPC; see the list above. The annotators annotated all data, not only the clear and/or obvious moments.

As a test run, a first interaction session (one pair of children interacting with one version of the system) was annotated by two annotators. Disagreements were identified and discussed, leading to a refinement of the original annotation instructions. Then one annotator used the revised instructions to annotate all interaction sessions. One of the interaction sessions (different from the first) was annotated by a second annotator, to establish inter-annotator agreement. For this we used Brennan and Prediger's free marginal kappa ("corrected kappa" in Anvil) because "Free-marginal versions of kappa are recommended when raters are not restricted in the number of cases that can be assigned to each category, which is often the case in the typical agreement study." (Randolph et al. 2005, p. 9). In our inter-annotator agreement study Brennan and Prediger's free marginal kappa turned out to be 0,5140, indicating low agreement. Although less suitable for our study, for comparison we also provide Cohen's kappa, which was 0,3389.

Clearly, annotating the level of social interaction is very subjective. One cause of the low agreement was that offering and accepting help did not always get recognised and annotated as cooperative play by the second annotator. Another problem was that the annotators labeled the interactions at different levels of granularity. For example, what was considered by one annotator as one long dialogue sequence (one segment labelled as highly social play) was annotated by the other as five separate segments of associative play interspersed with brief bursts of social play.

Given the low inter-annotator agreement and the small scale of the evaluation experiment we decided not to pursue a further quantitative analysis of the data. Instead we opted to use the carefully annotated videos as a basis for a more qualitative analysis, described in Section 4.3.

4.3 Observations

For the qualitative analysis of the children's interaction sessions, we used the SPC annotations to inspect those video segments that had been labeled as belonging to either the most social domains (highly social play and cooperative play) or the least social domain (different focus of attention). Based on our observations while reviewing these video segments, we tried to determine the reasons (or: 'triggers') for the displayed behaviour. Trying to find the reasons behind the children's behaviour can help us understand which system aspects to strengthen and which things to avoid in order to stimulate social interaction.

4.3.1 Highly social behaviour and cooperation

Inspection of those video segments where the children displayed highly social and cooperative behaviour revealed the following 'social triggers' in our setting.

Fun. One of the highly social behaviours that we frequently observed was combined eye contact and laughter (belonging in the highly social domain, according to the SPC). This behaviour was typically displayed after the occurrence of 'funny' actions, sounds and sentences in the developing story. In several cases the children were just laughing about the same thing at the same time, which is a more coincidental than intentional form of social interaction. So in that sense, this social behaviour could be seen as somewhat superficial. On the other hand, these moments of shared fun did create a relaxed atmosphere between the participants, and we observed that they were often more inclined towards further social interaction afterwards. A strong example of this are pair 3, two brothers who were quite shy and mainly communicated through shared glances and rare whispers. After they had shared laughter over Red performing a funny action ('fart under water' at the lake), they started talking together much more freely.

Interaction between the players' characters. Highly social behaviour between players often originated from social interaction between their characters. Examples of this behaviour are the player characters giving each other cakes, talking to each other in the story world, or becoming angry at each other. An example is the following exchange:

> **Pair 4, session 2**
> Action by Child 1: Red gives Grandma the apple pie.
> Child 2: *I am going to give it back; I baked it myself.*
> Action by Child 2: Grandma gives Red the apple pie.
> Child 2: *You have to eat it now.*
> Child 1: *No, you're trying to put something in it again.*

We find social interaction based on such in-story events quite valuable from a storytelling perspective, as we see it as the first step towards dramatic play.

Planning for the next actions or goal. Another kind of strong story-related social behaviour occurred when the children were making plans for the next actions of their character, or discussing a specific goal they had within the

story. This type of social interaction typically occurred late in the first interaction session or in the second, because initially the players were more occupied with discovering the interface and the story world. Only after some time they started to pay more attention to the story and the sequence of actions they performed. When they started to plan their actions in advance, rather than just selecting actions at random, this often resulted in telling the other player about their plans and goals. An example:

> **Pair 1, session 1**
> Action by Child 1: Grandma poisons the chocolate cake.
> Child 2: *I am going to tell Wolf that you took my chocolate cake!*
> Child 1: *No, I am going to poison it.*
> Child 2: *Yes, but then I'll give it to Wolf.*

Acquiring help to obtain story goals. At times, the children needed the help of the other player to obtain some of their short or long term story goals. This need for help then resulted in social interaction such as requesting the other player to perform a particular action or persuading him or her to cooperate towards a personal goal. This persuasion can result in players starting to have a shared goal for their story and to continue in cooperation. An example:

> **Pair 4, session 2**
> Action by Wolf: Wolf greets Red.
> Child 1: *You have to give him the cake!*
> Child 2: *It has to be poisoned first.*
> Child 1: *(...) I will do that, alright?*
> Child 2: *OK.*

Explaining character actions. After choosing an action for their character, players often explained the reason why they chose that particular action to the other player; see the following examples:

> **Pair 2, session 2**
> Action by Child 1: Red eats the birthday cake.
> Child 2: *Actually, Red should have given the birthday cake to Grandma.*
> Child 1: *I already ate it. Otherwise the wolf would have eaten it.*

> **Pair 3, session 2**
> Action by Child 1: Grandma bakes the cheese cake.
> Action by Child 2: Red bursts out in tears.
> Child 2: *Why can't I bake the cake?*

Also, when the children gained new insights about the story world or Wolf's behaviour, they often told the other player about the explanation they had come up with. We regard this type of highly social behaviour as very desirable because it contributes to a shared understanding and may contribute to the development of a more coherent story.

Waiting for Wolf and/or reflections on Wolf's behaviour. For reasons explained in Section 3.5, players often had to wait while Wolf was busy making a plan. On average, this took up around 27% of the total interaction time. These waiting periods were sometimes experienced as annoying, but they often also were an occasion for a social intermezzo between the players. During these intermezzos they reflected on Wolf's behaviour, or on what had happened so far and what they were planning to do in the next few rounds (for example, baking a cake to poison Wolf). An example of children discussing Wolf's plan making:

> **Pair 2, session 2**
> (Wolf is thinking)
> Child 1: *Wolf is trying to think of a plan.*
> Child 2: *Yes, but he can't do that because the cake has been eaten.*

Providing help. Helping each other was by far the most common reason for cooperative behaviour, the highest domain in the SPC. Often help was offered about things that were unclear about the system to one player but not the other, or that did not work as expected. An example:

> **Pair 1, session 2**
> (Child 1 is trying to move his character to a non-adjacent location)
> Child 2: *No, you can't do that.*
> Child 1: *Why can't we both be there?*
> Child 2: *You can't because you have to go that way. [points]*

Turn taking. Players occasionally reminded the other that it was their turn. On the surface, these reminders are cooperative behaviour, but in some cases they probably had a less social underlying motivation, i.e., the player wanting to reduce the time spent waiting for the other to perform their turn. Examples include, *Now it's your turn and then it's mine again* and *Hey, come on, hurry up! Otherwise I will do it.*

Bugs, flaws and quirks of the system. Occasionally, the children noticed some flaws or quirks in the system, for example delayed audio feedback, input recognition issues, or strange wording of the sentences used to report the characters' actions (*"Grandma bakes* the *apple pie", shouldn't that be "Grandma bakes* an *apple pie"?*). A positive side effect of these issues was that they triggered the players to talk about the system, the story, its interface and their expectations. Little deviations of what they had expected elicited players to express their thoughts and sometimes got a little dialogue going.

4.3.2 Different focus of attention

Different focus of attention is the least social of all annotation labels we assigned. In these cases, the children were not paying attention to each other's actions at all, which meant a total lack of social interaction. We observed a few clear reasons for this behaviour.

Waiting for Wolf. Although waiting for Wolf often seemed to work as a trigger for (highly) social interaction, it more frequently induced the participants to lose their shared focus of attention. When having to wait for Wolf, players often started looking around, playing with different interface elements or reading the story so far in their own story area.

Waiting for the other player. Because of the turn taking element of our board game-inspired interface, after each action the player had to wait for the other player to finish their turn before they could choose another action for their character. We found that during this waiting time, players usually paid close attention to what the other player was doing, but a few times having to wait for the other player resulted in losing the shared focus of attention.

'Just' distracted. On several occasions, the children lost their shared focus of attention because of various distractions, as there were a lot of interesting things to see in the experiment room. Besides looking around the room, it also happened a few times that a child's attention was drawn to the parent, the facilitator or one of the cameras in the room. In most of these cases it seemed that the children were not so much distracted by these people or objects, but that they served more as a target to focus on for the children while they were otherwise distracted. At times, one of the players just started looking at the floor or ceiling of the room, possibly because they were thinking about the story or the system.

Playing around with the story area. The story areas turned out to be among the most interesting elements in the interface, because they were resizable and could be scrolled. The children did not only use them to read the story, but they also used them as play objects, moving them around and trying to make them as big or small as possible. This playing sometimes resulted in losing the shared focus of attention. Sometimes one child was playing with the story area while the other was doing something else, but it also happened that they were both busy with their own story area without paying attention to the other.

4.4   Interview results

The informal interviews we held with the children were aimed at finding out their preference for one or the other version of the system (tangible or touch-only), their perception of social interaction when playing with either version, and their thoughts about the multi-player nature of the Interactive Storyteller. We were also interested in their previous experience with touch interaction.

Many young children in The Netherlands have experience in playing games on multi-touch devices such as smartphones or tablets. In fact, all children in our research answered that they often played games on smartphones or tablets at home or somewhere else. This experience might have influenced their expectations and ease of interaction with the interface, but we did not investigate this any further.

When asked about their preference for either of the two versions of the system, the children's preference for one version was always based on having trouble with the other version. Some children experienced problems when dragging the image of their character over the screen of the multi-touch table; they liked the

tangible version better. Others noticed that the system sometimes needed quite some time to recognise the identification marker underneath the tangible that represented their character; those players preferred the touch-only version. Based on our limited research it seems that it doesn't matter which version of the system players use, as long as the movement/recognition of characters is working properly and players do not experience any trouble with it.

The children's answers to the questions about in which version they had most often cooperated, talked, laughed, had eye contact, and had fun, did not show any differences in experienced social behaviour or fun between tangible and touch-only interaction. (In the annotation of the videos such differences could also not be found.) A remarkable result was that all eight children said they had laughed the most while using the second version, whichever version that was. In general it can be said that their answers depended almost exclusively on the order in which they had used the two versions, and appeared to be independent from the interaction style they had used. If they answered one of these "in which version most often" questions with the first version they had used, this was almost always because when playing with that version it was all still new to them and they were just getting to know the system, which resulted in more social interaction. If a child answered one of these questions with the second version, it was mostly because during the second interaction they had discovered new actions that triggered new ideas and resulted in one or several kinds of social interaction. In other words, their increased familiarity with the system gave them a sense of control and allowed them to create their own goals (two important elements of a 'flow' experience, according to Csikszentmihalyi (1991)), which in turn provided them with more grounds for social interactions such as communication and collaboration.

All children answered that they wanted to play with the system again. Some children answered that they would like to play with both versions again, others only mentioned the version with which they had experienced no trouble in moving their character. A few children answered that they wanted to play again with the version they had used in their second interaction session because it had more 'cool' actions, although the set of possible actions was exactly the same in both versions.

Most children answered that they would like to play with the system together rather than alone. The answers on the question about group size in possible larger groups surprised us in a positive way. Without being guided by the question, several children suggested that they wanted to play the game with their whole family (3-5 people). This result matches the social setting we initially had in mind when choosing our board game metaphor.

4.5 Example of a created story

To illustrate the kind of stories that were created in the evaluation experiment, Figure 6 shows a part of one of the stories created by the children, translated from Dutch to English. The sentences were generated using simple templates, reporting the events as they happened. An alternative option that is available in our storytelling framework, but was not used in our experiments, is to generate a full text after the story is finished. This makes it possible to use more advanced language generation techniques resulting in a more fluent text (Theune et al. 2007).

Once upon a time, there was a little girl with a red cap. She wanted to bring a birthday cake to her grandmother...

Wolf is mean. Red bursts out in laughter. Grandma bakes the cheese cake. Wolf sneaks to Red's house. "Hello, Wolf" says Red. With a little bit of cyanide, Grandma poisons the cheese cake. "Oh, hey, Red!" says Wolf. Red does the dishes. Grandma shuffles to the beach. "Give me the birthday cake", says Wolf and forcefully takes it away from Red. "Give that back!" says Red and takes the birthday cake back from Wolf. Grandma shuffles to Red's house. "Give that back!" says Wolf and takes the birthday cake back from Red. "Oh Grandma," says Red, "Wolf stole the birthday cake from me!" "Oh that is not nice, I feel very sorry for you" says Grandma to Red. Wolf eats the birthday cake. "Here you go, Wolf" says Grandma, giving the cheese cake to Wolf. Red bursts out laughing. Grandma bursts out in tears. Wolf becomes very hungry. "What should I do, Grandma?", asks Red. "Well, Red" says Grandma, "I have a plan to poison Wolf. You just wait and see!" Wolf eats the cheese cake.

Figure 6   Part of a story created during the experiments.

The actions chosen by the children sometimes may seem somewhat random; for instance, in the example below, Grandma bursting out in tears for no apparent reason. Sometimes such actions were just chosen for the fun of acting strange or unexpectedly (which was the case in the example below), but often the children's comments revealed story-related motivations. For instance, one child provided the following in-character explanation for their character bursting out in tears as the other child's character was moved to a different location: "Because you are leaving me!"

5 Discussion and future work

Our investigation which elements of the Interactive Storyteller encouraged or discouraged social interaction between the players, led to some interesting observations. Certain properties of the system that might seem negative at first sight, turned out to have unexpected benefits for social interaction.

5.1 Waiting times

Games that use turn taking inherently introduce waiting time for each of the players. In the case of the Interactive Storyteller, this waiting time was sometimes unwelcome for players who were eager to go on playing with their own character. It however also 'forced' players to pay attention to what the other players did, simply because they could not do much else at that time. This led to higher involvement in the actions the other player chose, and thus provided more options for social interaction. Interesting research in this field is currently being done by Pape and Graham, who are exploring how the turn-based game play of board games on digital tabletops can be relaxed towards more liberal coordination policies, in such a way that the social advantages of board games can be preserved (Pape & Graham 2010).

Another cause of waiting time for the players was the long periods it often took for Wolf to realise there was no plan to satisfy the adopted goal. Sometimes these long waiting times resulted in the players getting distracted while waiting, but in many cases they seemed to work as a trigger for social interaction in the form of story-related discussions. This was probably helped by the fact that while waiting for Wolf, both players had no other occupation, unlike in the case of turn taking where one was busy while the other was waiting. This suggests that even when real-time computing for system-controlled characters is technically feasible (for example, when using more efficient planning algorithms) it might be a good idea to build in some forced moments of reflection. During these periods none of the players can perform any actions and they are invited to discuss past and future story developments. By comparing an interaction condition with controlled pauses to one without delays, we could then experimentally test our hypothesis that social interaction benefits from shared waiting times.

5.2   Cooperation and goals

Cooperative behaviour mostly occurred when the players were helping each other with the interface, something that should no longer be necessary when the interface is further improved. However, cooperation also sometimes took the form of working towards a particular goal together, and this behaviour is a type of social play we would particularly like to stimulate in the Interactive Storyteller. One way to do this would be to increase the number of character actions that can globally affect the storyline.

The version of the LRRH story domain we used in our user evaluation includes many character actions that are local in nature, for example doing the dishes or diving into the lake. These actions provide the players with 'local agency' in the sense that they have immediate and visible effects on the story (Mateas & Stern 2005). However, they do not provide the player with global agency in the sense that they have no major influence on the course of the story. As a consequence, these actions are not useful for helping (or hindering) the other player to achieve some story goal. To support collaboration and competition, more actions are needed that are related by causal links, allowing them to function as building blocks for long-term plans. This is in line with the authoring approach advocated by Swartjes & Theune (2009) to achieve more coherent story domains.

In relation to goal-oriented cooperation, it should be noted that in our evaluation experiment we did not give the participants any specific goals. Still, we observed that as the children became more familiar with the possibilities of the system, they started setting themselves personal goals and challenges. These were often story-related (trying to poison Wolf) but sometimes they were also related to interface elements, for example trying to create the biggest possible story area. According to the influential flow theory of Csikszentmihalyi (1991), having clear goals is one of the important elements required to achieve 'flow', a state of optimal enjoyment. Other flow elements are immediate feedback on actions and a sense of control, which are both supplied by the Interactive Storyteller.

5.3 Tangibility

Our observations of the children's interaction sessions and the results of the informal interviews did not reveal any influence of interaction style (tangible or touch-only) on the players' social behaviour. The use of tangibles can be useful for many different applications (Hornecker & Buur 2006), but we did not observe any advantages for tangible playing pieces to represent player-controlled characters in our storytelling board game. This may well be due to the fact that in the tangible version of our system, we chose not to fully exploit the possibilities offered by 3D tangibility. To enable a fair comparison between tangibles+touch and touch-only in our experiment, we minimised the differences between the two versions, avoiding any functionality that could only be implemented in one of the two. However, for future versions of the Interactive Storyteller, we might consider designing a role for the tangible playing pieces that is greater than just a tangible version of graphical icons.

A disadvantage of the use of tangibles is that it imposes quite a few restrictions on the system and constrains the flexibility of the interface. The system has to deal with tangibles being placed anywhere on the map, including disallowed destinations or places that are not actually a location. Also, when tangibles are placed on particular places on the map, the system should never rotate, zoom or tilt the map, because it can not physically move the tangibles to their changed position. We therefore suggest that the next prototype of the Interactive Storyteller should not use tangible playing pieces, but should be based on touch-only user input.

Another advantage of using images to represent characters instead of tangibles is that it offers the possibility of customisation. Fantasy, in the form of imaginary characters with which the individual can identify, can be a strongly motivating factor to engage in some activity (Malone & Lepper 1987, Korhonen et al. 2009). By tailoring the look of the character to the user, identification can be increased. In theory, both tangible and graphical representations of characters could be customised, but a character image editor is much more practical and flexible than customising physical playing pieces. Introducing a simple character image editor like those available in many contemporary computer games is almost without effort and could provide in the need, expressed by some children, to make their character more "their own".

5.4 Future Work

One of the future research challenges to be addressed is allowing the possibility to add engaging visual elements (e.g., object inventories, animations) while keeping the interface framework generic for use with other story domains. Adding such visual elements would further strengthen the engagement factor that Malone & Lepper (1987) called 'sensory curiosity', which we found to be a strong trigger of social behaviour. On the other hand, this social triggering effect is likely to diminish over time as the novelty of the audiovisual effects wears off. Lasting influence on social interaction can probably be better achieved by methods that are more structurally related to the interactive storytelling, e.g., providing more possibilities for character interaction and player collaboration.

The scalability of the ASI for larger story domains with more choice should be further improved. Currently the possible actions in the ASI are only grouped by action type. However, if the character Red has four cakes and there are two characters she can give the cakes to, this already results in a list with eight possible 'give' actions. Larger domains with even more objects and characters can lead to an explosion of possibilities. In this case, the number of options to choose from in the ASI should be reduced by first asking the user to select the patient of the action before displaying the list of (remaining) possible instantiations of the action. If this still leaves many possibilities, this step can be repeated for the 'recipient', 'instrument' and other applicable thematic roles. The possible arguments of every action predicate are already available in our storytelling framework. Larger story domains also require a larger map and more graphical elements to represent the characters and locations. The extent to which this is possible mainly depends on the hardware used to display the story world. The scalability of the map of any domain is limited by screen size, screen resolution, finger size, finger recognition and the minimum marker size for tangible recognition.

Finally, in the future we would like to port the Interactive Storyteller to the new generation tablet computers. These devices live up to the high expectations and increasing demands of present-day users and provide more accuracy for new directions in multi-touch research. Several of the participants in our evaluation experiment expressed an interest in using the Interactive Storyteller in a family setting. One way of bringing the Interactive Storyteller to the living room might be to play it on a large tablet lying on the table like the board of a traditional board game, which would bring us even closer to our original vision of the Interactive Storyteller as a 'board game with benefits'.

Acknowledgements

References

Alves, A., Lopes, R., Matos, P., Velho, L. & Silva, D. (2010), Reactoon: Storytelling in a tangible environment, *in* '3rd IEEE International Conference on Digital Game and Intelligent Toy Enhanced Learning', pp. 161–165.

Aylett, R., Louchart, S., Dias, J., Paiva, A. & Vala, M. (2005), FearNot! - an experiment in emergent narrative, *in* 'Intelligent Virtual Agents', Vol. 3661 of *Springer LNCS*, pp. 305–316.

Broadhead, P. (2004), *Early years play and learning: developing social skills and cooperation*, Routledge.

Broadhead, P. (2006), 'Developing an understanding of young children's learning through play: the place of observation, interaction and reflection', *British Educational Research Journal* **32**(02), 191–207.

Broadhead, P. (2009), 'Conflict resolution and children's behaviour: observing and understanding social and cooperative play in early years educational settings', *Early Years* **29**(2), 105–118.

Cappelletti, A., Gelmini, G., Pianesi, F., Rossi, F. & Zancanaro, M. (2004), Enforcing cooperative storytelling: First studies, *in* '4th IEEE International Conference on Advanced Learning Technologies', pp. 281–285.

Choi, D. & Kim, J. (2004), 'Why people continue to play online games: In search of critical design factors to increase customer loyalty to online contents', *CyberPsychology & Behavior* **7**(1), 11–24.

Coldewey, D. (2010), 'Hands-on: D&D on the Microsoft Surface', http://techcrunch.com/2010/02/10/hands-on-dd-on-the-microsoft-surface/. Last accessed 13 April 2013.

Csikszentmihalyi, M. (1991), *Flow: The Psychology of Optimal Experience*, HarperCollins.

Dietz, P. & Leigh, D. (2001), Diamondtouch: a multi-user touch technology, *in* 'Proceedings of the 14th annual ACM symposium on User interface software and technology', ACM, pp. 219–226.

Helmes, J., Cao, X., Lindley, S. & Sellen, A. (2009), Developing the story: Designing an interactive storytelling application, *in* 'ACM International Conference on Interactive Tabletops and Surfaces', pp. 49–52.

Hornecker, E. & Buur, J. (2006), Getting a grip on tangible interaction: A framework on physical space and social interaction, *in* 'Proceedings of CHI 2006: Designing for Tangible Interactions', pp. 437–446.

Karadimitriou, K. & Roussou, M. (2011), Studying player experience in a collaborative embodied interaction game, *in* 'Third International Conference on Games and Virtual Worlds for Serious Applications (VS-GAMES)', pp. 199–206.

Kipp, M. (2001), Anvil – a generic annotation tool for multimodal dialogue, *in* 'Seventh European Conference on Speech Communication and Technology', pp. 1367–1370.

Korhonen, H., Montola, M. & Arrasvuori, J. (2009), Understanding playful experiences through digital games, *in* 'Proceedings of the 4th international Conference on Designing Pleasurable Products and Interfaces (DPPI '09)', pp. 274–285.

Kriegel, M., Lim, M., Aylett, R., Leichtenstern, K., Hall, L. & Rizzo, P. (2010), A case study in multimodal interaction design for autonomous game characters, *in* '3rd Workshop on Multimodal Output Generation', pp. 15–25.

Kruger, R., Carpendale, S., Scott, S. & Greenberg, S. (2003), How people use orientation on tables: Comprehension, coordination and communication, *in* 'Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work', pp. 369–378.

Magerkurth, C., Memisoglu, M., Engelke, T. & Streitz, N. (2004), Towards the next generation of tabletop gaming experiences, *in* 'Proceedings of Graphics Interface 2004', pp. 73–80.

Magerkurth, C., Stenzel, R. & Prante, T. (2003), STARS – a ubiquitous computing platform for computer augmented tabletop games, *in* 'Proceedings of Video Track of Ubiquitous Computing (UBICOMP03)'.

Malone, T. & Lepper, M. (1987), Making learning fun: A taxonomy of intrinsic motivations for learning, *in* R. Snow & M. Farr, eds, 'Aptitude learning and instruction 3', Lawrence Erlbaum Associates, pp. 223–25.

Mandryk, R. & Maranan, D. (2002), False Prophets: Exploring hybrid board/video games, *in* 'Conference on Human Factors in Computing Systems', pp. 640–641.

Mateas, M. & Stern, A. (2003), Façade: An experiment in building a fully-realized interactive drama, *in* 'Game Developers Conference: Game Design Track'.

Mateas, M. & Stern, A. (2005), Structuring content in the Façade interactive drama architecture, *in* 'Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE '05)', AAAI Press, Menlo Park, CA, pp. 93–98.

Pape, J. & Graham, T. (2010), Coordination policies for tabletop gaming, *in* 'Proceedings of Graphics Interface 2010', pp. 24–25.

Piper, A., O'Brien, E., Morris, M. & Winograd, T. (2006), SIDES: A cooperative tabletop computer game for social skills development, *in* '20th Anniversary Conference on Computer Supported Cooperative Work', pp. 1–10.

Pizzi, D. & Cavazza, M. (2007), Affective storytelling based on characters' feelings, *in* 'Intelligent Narrative Technologies: Papers from the AAAI Fall Symposium', pp. 111–118.

Prada, R., Paiva, A., Machado, I. & Gouveia, C. (2002), "You cannot use my broom! I'm the witch, you're the prince": Collaboration in a virtual dramatic game, *in* 'Intelligent Tutoring Systems', Vol. 2363 of *Springer LNCS*, pp. 913–922.

Randolph, J., Thanks, A., Bednarik, R. & Myller, N. (2005), Free-marginal multirater kappa (multirater κfree): an alternative to fleiss fixed-marginal multirater kappa, *in* 'Joensuu Learning and Instruction Symposium', pp. 1–20.

Rubin, K. (1982), 'Play Observation Scale', *Child Development* **53**(3), 651–657.

Russell, S. & Norvig, P. (1995), *Artificial intelligence: A modern approach*, Prentice Hall.

Swartjes, I. (2010), Whose Story Is It Anyway? How Improv Informs Agency and Authorship of Emergent Narrative, PhD thesis, University of Twente.

Swartjes, I. & Theune, M. (2008), The Virtual Storyteller: Story generation by simulation, *in* '20th Belgian-Netherlands Conference on Artificial Intelligence', pp. 257–265.

Swartjes, I. & Theune, M. (2009), Iterative authoring using story generation feedback: Debugging or co-creation?, *in* '2nd International Conference on Interactive Digital Storytelling', Vol. 5915 of *Springer LNCS*, pp. 62–73.

Theune, M., Slabbers, N. & Hielkema, F. (2007), The Narrator: NLG for digital storytelling, *in* 'Proceedings of the Eleventh European Workshop on Natural Language Generation', pp. 109–112.

Tse, E., Greenberg, S., Shen, C. & Forlines, C. (2007), 'Multimodal multiplayer tabletop gaming', *Computers in Entertainment*. Article 12.

Tse, E., Shen, C., Greenberg, S. & Forlines, C. (2006), Enabling interaction with single user applications through speech and gestures on a multi-user tabletop, *in* 'Proceedings of the working conference on Advanced visual interfaces', ACM, pp. 336–343.