

UNIVERSITY OF TWENTE.

MASTER'S THESIS

The Interactive Storyteller

A multi-user tabletop board game interface to support social interaction in AI-based interactive storytelling

Thijs Alofs

August 16, 2012



Human Media Interaction
Department of Electrical Engineering,
Mathematics & Computer Science
University of Twente

Graduation Committee:
Mariët Theune
Betsy van Dijk
Frans van der Sluis

Abstract

This thesis presents *the Interactive Storyteller*, a ‘story domain’-independent multi-user interface for AI-based interactive storytelling. Up until now, users of our storytelling framework, which uses the *emergent narrative* approach, had no influence on how the storyline developed. In the current research we aim to offer users a social setting for AI-based interactive storytelling. Existing AI-based interactive storytelling systems do not offer its users a shared interface in a social setting.

To support face-to-face contact and social interaction, we position users around a shared multi-touch table. This very much resembles the social setting of traditional tabletop board games. We analysed video recordings of interactions that pairs of children had with our system to determine what triggered cooperation and highly social behaviour and what caused players to have a different focus of attention. We also tried to find out whether the use of tangible playing pieces to reinforce the resemblance with traditional board games offered any advantages over touch-only interaction, but we could not find any differences.

To our knowledge, the Interactive Storyteller is the first AI-based interactive storytelling system that combines the emergent narrative approach with the social aspects of traditional tabletop board games.

Preface

Contents

1	Introduction	1
1.1	The Virtual Storyteller framework	1
1.2	Research goal: The Interactive Storyteller	2
1.3	Story domain and target user group	3
1.4	Thesis outline	4
2	Related work	5
2.1	Interactive storytelling systems	5
2.1.1	Façade	5
2.1.2	FearNot!	6
2.1.3	EmoEmma	7
2.1.4	Teatrix	7
2.1.5	ORIENT	8
2.1.6	Conclusions	9
2.2	Tabletop storytelling systems	11
2.2.1	StoryTable	11
2.2.2	TellTable	11
2.2.3	Reactoon	12
2.2.4	Tangible Spatial Narratives (TViews)	12
2.2.5	Conclusions	13
2.3	Digital tabletop board games	15
2.3.1	STARS Monopoly and KnightMage	15
2.3.2	Warcraft III and The Sims	16
2.3.3	SurfaceScapes Dungeons & Dragons	16
2.3.4	Conclusions	16
3	Design	18
3.1	Requirements	18
3.2	Interface design	19
3.3	Early user tests	23
3.4	Pilot user test	24
3.5	Reflection on requirements	26
4	Evaluation	28
4.1	Experiment setup and procedure	28
4.2	Annotation method	31
4.3	Agreement analysis	34
4.4	Results of annotation and observations	35

Contents

4.5	Interview results	39
4.6	Example of created story	41
5	Discussion	43
5.1	Social interaction imposed by turn-taking	43
5.2	Tangibles	44
5.3	Cooperation and goals	44
6	Conclusions	46
7	Recommendations	48
7.1	New experiments	48
7.2	Customisation	49
7.3	Scalability	49
7.4	Visual feedback	50
7.5	Moments of reflection	50
7.6	Tablet board game	50
	Glossary	51
	Acronyms	53
	Bibliography	54
A	Original interview questions	61
B	Quantitative results	62

Chapter 1

Introduction

This thesis describes the design, implementation and evaluation of “The Interactive Storyteller”, the interactive branch of the Virtual Storyteller, a multi-*agent* story generation system. Up until now, users of the Virtual Storyteller had no influence on the story creation process and were mere spectators. User interaction has been a goal since the beginning of the Virtual Storyteller project [Theune et al., 2002].

Below, we first give a brief overview of the Virtual Storyteller framework. After that, section 1.2 introduces the idea behind the Interactive Storyteller and the goal of our research. We discuss the chosen story domain and target user group in section 1.3. Finally, in section 1.4 the further outline of this thesis is presented.

NOTE: glossary terms are displayed in italics on first use. Both glossary terms and acronyms are clickable in the digital version of this document and lead to the glossary and acronym list at the end of this thesis.

1.1 The Virtual Storyteller framework

The Virtual Storyteller¹ storytelling framework of the Human Media Interaction (HMI) research group is a multi-agent system that generates stories in a virtual story world [Swartjes & Theune, 2008]. In the Virtual Storyteller there is an intelligent agent for every character in a story that acts out the role of that character. Each agent makes and executes a plan consisting of a sequence of actions to satisfy its goals. The set of actions a character can perform at any particular moment is limited by the definition of possible actions in the story domain, the current state of the story world, and the mental state of the character. From the ongoing interaction between the agents and through their choices of action a story emerges; this approach in *AI-based* storytelling is called ‘*emergent narrative*’ [Aylett, 1999].

After the story generation process of the Virtual Storyteller is finished, a narrator component of the system can analyse the formal representation of the events in the story and transform this information into a fluent story [Theune et al., 2007]. It is also possible to present this well-formed narrative to users by a virtual presenter like DEIRA [Knoppel et al., 2008].

¹<http://wwwhome.cs.utwente.nl/~theune/VS/>

Up until the introduction of the Interactive Storyteller, there was no user interaction and no visual representation of the story world.

1.2 Research goal: The Interactive Storyteller

In the Virtual Storyteller, the story generation process is a combined effort of several autonomous agents that act as characters in the simulated story world. The users or ‘audience’ cannot give input or feedback and therefore have no part in how the story develops while it emerges.

The interface for the interactive version of the Virtual Storyteller should be domain-independent, as it should easily be used for any existing or future story domain. We decided to go for multi-user interaction in our interactive storytelling interface, because enjoyment in games often comes from interaction with other people [Lazzaro, 2008]. Also, people that play in groups have more and intenser emotions than when they are playing alone [Lazzaro, 2004]. Storytelling is, in its origin, a highly social conversational activity and therefore, social interaction is an important aspect of the storytelling experience.

Social interaction is an essential component in many models of what it is that constitutes an enjoyable, playful experience [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].

Social interaction can be defined as being “*the process by which people act and react in relation to others*” [Macionis & Gerber, 2010, p. 127]. This description is however quite abstract and not very suited to explain what we mean by ‘social interaction’ in this thesis. The following is a description that is much more appropriate to illustrate what kind of behaviour we want to support with our system:

“Social interaction is the use of nonverbal or verbal behavior to engage in interaction with people. This can involve eye gaze, speech, gestures, and facial expressions used to initiate and respond to interactions with others.”

[Wetherby & Wiseman, 2007, p. 21]

The **goal** of the current research is to design and evaluate a multi-user interface for an interactive storytelling system that is domain-independent and supports social interaction. Existing AI-based interactive storytelling systems do not offer a shared interface in a social setting to its users.

To support face-to-face contact and social interaction in computer games, Magerkurth et al. suggest to use “interaction metaphors that adhere to well-proven physical interfaces found in strongly human centred entertainment genres such as tabletop or board games” [Magerkurth et al., 2004, p. 164]. By providing users with a shared visual representation of the story world on a *multi-touch* table, we try to foster a shared focus of attention during the interaction with the system. With this shared focus of attention and the possibility of face-to-face contact, we think we have a social and intuitive approach for group interaction with the Interactive Storyteller. Imagine multiple users, all standing

or sitting around a multi-touch table together and playing some kind of tabletop story game. Each user can play a character in the story, but it is also possible that users collaboratively control the characters, discussing decisions for the characters with each other. A digital narrator can tell the story as it evolves step by step and/or afterwards in one fluent story without interruptions.

The resemblance with board games is the main metaphor we adopted for the interaction design of the Interactive Storyteller. Our storytelling framework follows a turn-based time line [Swartjes & Theune, 2008], which corresponds with turn-taking in many tabletop board games.

To reinforce the resemblance with traditional board games, we investigate the use of *tangible* playing pieces that represent characters for physical interaction. 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). We compare a version of the Interactive Storyteller that uses these tangible playing pieces and a touch-only version of the system to see whether the use of tangibles has advantages over a touch-only approach.

1.3 Story domain and target user group

Although we want the interface of the Interactive Storyteller to be domain-independent, we had to choose a story domain for the user experiments and the examples in this thesis. Two story domains are currently available in our storytelling framework. A domain about pirates, named PLUNDER, and a domain based on the “Little Red Riding Hood” story [Swartjes & Theune, 2009]. We used the latter domain in the experiments and in this thesis, because we consider it to be more coherent, easier to understand for new users, and easier to visualise than a pirate ship. Creating a new story domain was not an option because it requires a decent amount of authoring, for which no time was available.

Given that Little Red Riding Hood (LRRH) is a children’s fairytale and because children are an interesting and often used target group for interactive storytelling, our research is focussed on children. In The Netherlands, children start to learn reading and writing around the age of 6 and continue to follow primary education until they are at least 11 years of age. Our experience is that at the age of 8, most children are very able to quickly and easily read sentences (as used in our interface, see subsection 3.2.3) and simple stories, so our research is focussed predominantly on children aged 8 to 11.

So, in the current research we are interested in social interaction of playing children. We look at children’s play behaviour by using several categories with increasing levels of momentum and reciprocity in observed actions and language, in accordance with [Broadhead, 2009]. The category *Social play* is indicated by smiling, laughter, play noises, play voices, mutual eye contact, and isolated verbal interactions. The category *Highly social play* is indicated by dialogue, reciprocal sequences, clustered social play (eye contact with laughter), and development of joint play themes. The category *Cooperative play*, being the highest form of social behaviour, is indicated by offering and accepting help, problem solving, role play, and achieving shared goals.

1.4 Thesis outline

The remainder of this thesis is laid out as follows. Chapter 2 contains the results of a literature study that looked at comparable or otherwise related systems and what can be learnt from them. In chapter 3 we describe the interface design of the Virtual Storyteller and its underlying considerations. In chapter 4 we present our observations from a small-scale user evaluation, focusing on the question which factors stimulated or hindered social interaction. In chapter 5 the results of the user tests are discussed, conclusions are drawn in chapter 6 and recommendations for future work are presented in chapter 7.

Chapter 2

Related work

In this chapter we look at existing interactive storytelling systems and other fields of related research. We start with an overview of some well-known interactive storytelling systems in section 2.1 to introduce the broad field of interactive storytelling. After that, we look at some existing tabletop storytelling systems in section 2.2. The third group of related research is discussed in section 2.3 and consists of several digital tabletop board game systems. Each section concludes with an discussion on how the described systems relate to the Interactive Storyteller and what the differences and similarities are.

2.1 Interactive storytelling systems

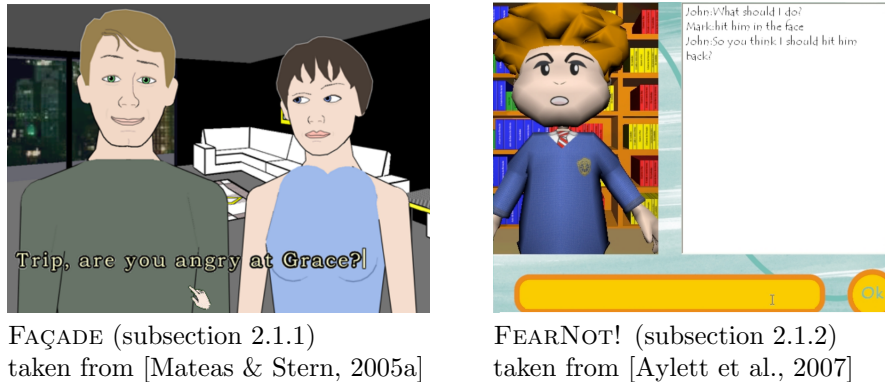
Although interactive storytelling is a relatively new research topic, already a lot of research has been performed in this field. Many experiments with possible approaches and different interfaces for interaction have been carried out.

The goal of this section is not to provide a complete overview of all interactive storytelling systems, but to present a selection of some of the most well-known, advanced, extended, similar or otherwise notable systems.

2.1.1 Façade

FAÇADE (see Figure 2.1 left) is widely recognised as being the first fully realised and published AI-based interactive digital storytelling system, or ‘interactive drama’ as the creators call it [Mateas & Stern, 2003]. In an interface similar to many real-time computer games, the user plays the role of a longtime friend of a couple whose marriage seems to be falling apart. The user visits this couple in their appartement where the user can walk around, engage in a conversation by typing text and use a hand cursor to point, gesture, touch or grab things.

The dynamically evolving storyline is not structured by obvious branch points, so the plot does not feel predefined or scripted. The unconstrained continuous natural language input resulted in high user expectations. Because of this free text input, choices are not limited to a predefined list of sentences. The system fails to understand approximately 30% of the typed input, but it avoids the “I don’t understand” response.



FAÇADE (subsection 2.1.1)
taken from [Mateas & Stern, 2005a]

FEARNOT! (subsection 2.1.2)
taken from [Aylett et al., 2007]

Figure 2.1: FAÇADE and FEARNOT!

FAÇADE has no educational goals for its users. It is all about (interactive) entertainment and providing an example of what is possible in the field of interactive digital storytelling by developing a fully realised system. In the game there are also no predetermined goals; users are free to decide what to do. To increase dramatic effect (and the user's influence on the storyline), the user is allured or even forced to choose sides in the couples' argument.

Finite state machines and scripting languages were considered ineffective to simulate complex character behaviour while handling continuous user input at the same time. An AI-based approach was adopted and 'procedural authoring' [Mateas & Stern, 2005a] was introduced to accomplish this goal. FAÇADE generates sequences of predefined sentences in response to free-text user input. Authoring these sentences took about three person-years alone [Mateas & Stern, 2005a].

Because of the many different possible outcomes, only after 6 or 7 sessions of about 20 minutes users feel they have exhausted all available options in FAÇADE. Anybody can try it, because FAÇADE is free for download¹.

2.1.2 FearNot!

FEARNOT!² (see Figure 2.1 right) is a narrative-centred learning environment (NLE) in which virtual drama is used as anti-bullying education [Aylett et al., 2005]. The system displays scenes which are inhabited by characters that are common in bullying situations (bullies, victims, assistants, bystanders, and defenders) in an interface that is quite similar to many computer games [Watson et al., 2007]. Each character is controlled by an autonomous intelligent agent and has a 3D cartoon-like embodiment.

Between scenes, the character that is the victim of the bullying asks the child user (9-11 years old) what to do next [Aylett et al., 2007]. They discuss a bullying coping strategy through a free-text interface. The advice given by the child user influences the mental state of the victim which often results in applying the proposed coping strategy in the next bullying scene. Because many

¹<http://www.interactivestory.net/>

²Fun with Empathic Agents Reaching Novel Outcomes in Teaching

different kinds of advice can be given, emergent narrative is used in the scenes instead of *scripted narrative*. FEARNOT! allows children to explore various coping strategies in bullying situations without this having consequences for themselves.

2.1.3 EmoEmma



Figure 2.2: EMOEMMA (taken from [Charles et al., 2009])

EMOEMMA (see Figure 2.2) is a AI-based interactive storytelling system based on the formalisation of the famous novel “Madame Bovary” by Gustave Flaubert [Pizzi & Cavazza, 2007]. The researcher’s abstract formalisation of several chapters of the second part of the novel was largely based on Flaubert’s own formalisation in which the writer describes the feelings and desires of the characters (Emma Bovary, Charles Bovary and Rodolphe Boulanger) in great detail. The availability of the formalisation of the story by the original author is probably the most important reason for choosing this story domain.

In the purpose-built multi-agent system, each character is independently controlled by a *heuristic search planner* in a shared world. The user can interact with the system by typing comments in natural language to one of the characters to influence their emotional state and beliefs. The Unreal TournamentTM engine is used for 3D visualisation, resulting in an interface similar to many 3D computer games.

2.1.4 Teatrix

TEATRIX (see Figure 2.3) is a NLE just as FEARNOT!. It is a 3D environment for AI-based collaborative story creation [Prada et al., 2000]. It was based on an existing classroom game for young children (7-9 years old) and combines acting, reading and writing with the goal to increase their narrative competence [Machado et al., 2000].



Figure 2.3: TEATRIX (taken from [Prada et al., 2002])

By choosing and combining scenes, characters and items, a story setup is created which defines the initial state of the virtual world [Prada et al., 2002]). After this initial phase, every child can select a character it wants to play during the story. Characters that have not been chosen by a child, will be initialised as autonomous agents. During the story creation phase, children can control their character by choosing an action from a defined set. A story director agent is present in the background to guide the story and to ensure coherence and believability. Afterwards, the created story can be viewed back as a movie and commented on with text by the children.

The children that use TEATRIX are sitting in the same classroom and can talk to each other, but they all use their own computer for interaction with the system. The view users have on the story world automatically follows the character they control, similar to some 3D computer games.

2.1.5 ORIENT

ORIENT¹ (see Figure 2.4) is an AI-based inter-cultural collaborative role-playing game (RPG) [Kriegel et al., 2008]. It uses the emergent narrative approach to generate the storyline. There is no predetermined plot and according to Kriegel et al. [2008] most authoring work was related to configuring character goals, actions, emotions and motivations. The system tries to improve the integration of children of different cultural backgrounds by increasing inter-cultural empathy of users [Kriegel et al., 2010]. In that light ORIENT can be seen as a *serious game* or narrative-centred learning environment (NLE). This feeling of understanding should be obtained through empathic behaviour in a role-playing game.

A team of three teenage users lands on a small alien planet and have to become friends with its virtual inhabitants, the Sprytes. These humanoid lizard-like game characters are controlled by autonomous affective agents. The human team shares a first person view on the 3D virtual world which is projected on

¹Overcoming Refugee Integration with Empathic Novel Technology

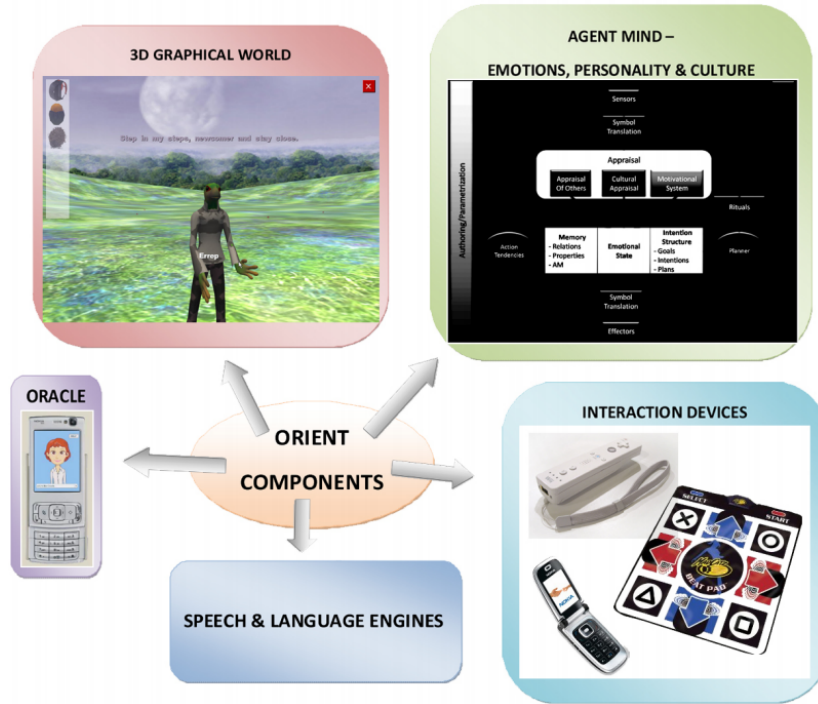


Figure 2.4: ORIENT (taken from [Kriegel et al., 2010])

a large screen. There are just four locations on the planet and between them only predefined pathways can be followed.

Each user in the team has their own device to interact with the world and the virtual characters in it. The user responsible for navigation and movement through the world controls a Dance Mat with his or her feet. For nonverbal communication with the Sprytes another team member uses a WiiMote that can recognise gestures. A mobile phone with speech recognition is used by the third user to name the Spryte to which the gesture with the WiiMote is addressed. This phone also has an RFID scanner to link tangibles in front of the users to virtual objects. On a second mobile phone a 2D tutor-like character is available to explain situations or Spryte behaviour to the users in order to enhance their learning of the Spryte culture and to prevent stagnation in the development of the storyline.

2.1.6 Conclusions

Although interactive storytelling is a relatively young research field, already many experiments with possible interfaces for it have been performed. We have looked at several interactive storytelling systems, all of which include some kind of intelligence by using agents or other AI technology. 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. TEATRIX and ORIENT focus

on collaborative storytelling by multiple users, but these also have interfaces similar to computer games. The interface of the Interactive Storyteller is very different from other AI-based storytelling systems, because it offers a multi-user interface on a shared tabletop surface. In section 2.2 and section 2.3 we look at systems that do match our tabletop platform, but are less similar to our AI-based storytelling framework.

In this section we have seen that the use of interactive storytelling systems is often a serious game or narrative-centred learning environment (NLE) for training and education of children, for example FEARNOT!, TEATRIX and ORIENT. FAÇADE and EMOEMMA in contrast are more about entertainment, art and providing a glimpse of what is possible in the field of interactive digital storytelling. For now, the Interactive Storyteller belongs to the latter category, but in the future a new domain for the Interactive Storyteller could be developed with the goal to train or educate children or adults in a serious game or NLE.

Just like our own storytelling framework, ORIENT and FEARNOT! use the emergent narrative approach to generate the storyline. Although all AI-based interactive storytelling systems are not completely scripted like *branched narrative* computer games, still much authoring has to be done in advance. Therefore, there are only a few fully realised systems and most systems have very limited domains. A fully realised system such as FAÇADE required several years to develop, although it only provides interactions of 20 minutes and is exhausted after a few times playing. Because our primary focus is on the tabletop interface, for the Interactive Storyteller we keep the authoring part to a minimum. For now, we focus on developing a generic domain-independent interface and not on extending the existing LRRH domain by investing a lot of authoring time.

2.2 Tabletop storytelling systems

We have seen some general examples of interactive storytelling systems in the previous section. As mentioned in the introduction of this thesis, we suggest to use a tabletop interface for our interactive storytelling system. The idea of using tabletop interfaces for digital storytelling is however not entirely new. In this section we will look at several storytelling systems that present their users a tabletop interface for interacting with the system.

2.2.1 StoryTable



Figure 2.5: STORYTABLE (taken from [Gal et al., 2009])

STORYTABLE (see Figure 2.5) supports cooperative storytelling for groups of children [Cappelletti et al., 2004]. The virtual story world is populated by ladybugs that represent the tools to create the story. The largest ladybug can be used to select a background from a collection to change the setting for the story. Other ladybugs, with the same but smaller shape, carry objects and characters that can be used in the story. The children can use a third type of ladybugs with a different shape to record audio snippets that will form the story. Objects and characters can be moved around, but these movements are not recorded by the system for playback. The STORYTABLE system uses the MERL DiamondTouch tabletop. This multi-touch system can recognise simultaneous touches of different users, but can not recognise tangibles. Important decisions are synchronised by using an explicit cooperative multi-user action, like everybody touching a button at the same time. The StoryTable system was used for reflective storytelling after a museum visit [Zancanaro et al., 2007] and the system was used to enhance social skills in children with High Functioning Autism [Gal et al., 2009].

2.2.2 TellTable

TELLTABLE (see Figure 2.6 left) was developed at the Microsoft research centre in Cambridge for the Microsoft Surface multitouch table [Helmès et al., 2009]. It can be used by children to create 2D images of objects and characters and facilitates the recording of stories. TELLTABLE aims to stimulate creativity and collaboration in storytelling for children. The creation of characters, setting and



TELLTABLE (subsection 2.2.2)
taken from [Helmes et al., 2009]



REACTOON (subsection 2.2.3)
taken from [Alves et al., 2010]

Figure 2.6: TELLTABLE and REACTOON

story are completely free. There are photography and drawing tools available to enable children to prepare their own characters and story environment before recording a story. While recording and narrating the story, the visual elements in the story can be directly changed or moved to resemble the story told.

2.2.3 Reactoon

The idea behind REACTOON (see Figure 2.6 right) is to support storytelling by letting children (aged 5 - 9) collaborate in creating a 2D animation for their story [Alves et al., 2010]. Creating and sharing stories is supposed to train and improve their communication skills. The system was designed with the use of tangibles in order to close the gap between a physical toy box and digital-only systems. Users are not required to have reading or writing skills, the creation of animations is achieved by using a combination of tangibles and multi-touch interaction. A story setting can be selected and characters and objects can be added to the story world. A tangible with a camera symbol is used to record scenes. To add additional sounds to a recorded scene, a tangible with a microphone symbol is available. The scenes with corresponding animations and sounds can be replayed one by one to view the complete story.

2.2.4 Tangible Spatial Narratives (TViews)

Tangible Spatial Narratives (see Figure 2.7) is a system for exploring a large existing multimedia story [Mazalek et al., 2007]. It uses the TVIEWS tabletop platform for tangible user interaction. Each tangible (pawn or puck) uses an electromagnetic sensor to determine its position and infrared transmission for data communication to the table. A top-projected map of the story environment provides a spatial framework for the story and the user interactions on the table [Mazalek & Davenport, 2003]. Multimedia clips are structured based on location, character and time. Pawns that represent characters in the story can be placed on the map and this reveals thumbnails of content clips corresponding to the selected character, location and time (controlled by using a clock tool). By selecting a thumbnail with a selection tool, the associated content clip is displayed on an additional upright screen with audio speakers. After viewing




Figure 2.7: TVIEWS tabletop platform, used for Tangible Spatial Narratives (taken from [Mazalek & Davenport, 2003])

the clip, the narrative advances and new content becomes available for playback. Step by step more is revealed and users move forward in their exploration of the story space. By navigating through the large collection of perspectives users experience an interactive narrative in this tabletop storytelling system.

2.2.5 Conclusions

The first three systems discussed in this section mainly focus on facilitating storytelling and supporting children in their development of narrative competence. These systems try to stimulate collaboration and creativity by offering a multi-user shared tabletop interface. None of them are a NLE or serious game with a particular goal, but more about developing skills for creating/telling stories and dramatic play in general. Social interaction during dramatic play is also an important aspect for the Interactive Storyteller.

Both TELLTABLE [Helmes et al., 2009] and REACTOON [Alves et al., 2010] can record animations and use tangible playing pieces for additional interaction. STORYTABLE does not use tangibles and can only record audio snippets [Cappelletti et al., 2004]. One of the two versions of the Interactive Storyteller also uses tangibles, but unlike above mentioned storytelling systems, it does not record animations or audio snippets. Our system is also less free in creating any kind of story. Aside from the use of a tabletop interface, the Interactive Storyteller is very unlike the storytelling systems discussed in this section and more like the AI-based systems discussed in section 2.1.

The system TANGIBLE SPATIAL NARRATIVES can be considered the odd one out in this section. It lets users move pawns around on a tabletop surface to explore an interactive story [Mazalek & Davenport, 2003]. Because this storytelling system is structured around an already existing story, it is kind of different from our system and the other tabletop storytelling systems discussed in this section. Just like our storytelling system however, the users of TANGIBLE SPATIAL NARRATIVES can explore an existing story world by moving character pawns around and make choices from a limited set of options.

Unlike our own storytelling system, the storytelling systems discussed in this section do not use system-controlled characters or other forms of AI to contribute to the story. In section 2.3 we will look at some systems that use some artificial intelligence in combination with the social advantages of tabletop systems.

2.3 Digital tabletop board games

After looking at some generic interactive storytelling systems and some tabletop storytelling systems in the previous sections, we will now briefly scratch the surface of the field of digital tabletop board games.

Unlike interactive storytelling systems, digital tabletop board games are not explicitly focussed on storytelling aspects. However, because many computer games develop a storyline, they are somehow (implicitly) related to interactive storytelling. Digital tabletop board games try to combine the dynamics and intelligence of computer games with the social advantages of traditional board games, just like we want to achieve with the Interactive Storyteller.

The system FALSE PROPHETS [Mandryk & Maranan, 2002] was one of the first to explore hybrid video/board games. In the following subsections several interesting and/or more recent digital tabletop board games will be described.

2.3.1 STARS Monopoly and KnightMage



MONOPOLY



KNIGHTMAGE

Figure 2.8: Games on STARS platform (taken from [Magerkurth et al., 2004])

STARS¹ is a platform for developing computer augmented tabletop games (see Figure 2.8) [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 audible 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, for instance the role playing game KnightMage and an adaptation of Monopoly [Magerkurth et al., 2004].



Figure 2.9: WARCRAFT III (left) and THE SIMS (right) on the DiamondTouch table (taken from [Tse et al., 2007])

2.3.2 Warcraft III and The Sims

The paper [Tse et al., 2007] demonstrates the application of a multi-user, multi-touch interface on existing single player computer games (see Figure 2.9). For this the DiamondTouch hardware from Mitsubishi Electric Research Laboratories (MERL) [Dietz & Leigh, 2001] is used. Input from the multi-touch table and speech has been mapped to the games WARCRAFT III and THE SIMS with the Gesture Speech Infrastructure [Tse et al., 2006]. The research shows that a social setting for game interaction can be provided by a tabletop interface. 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. Both games can be seen as an emerging interactive story, that (in theory) can be saved and told by a narrator.

2.3.3 SurfaceScapes Dungeons & Dragons

The SurfaceScapes team¹ (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 (see Figure 2.10) [Coldewey, 2010]. A very 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.

2.3.4 Conclusions

Just like we try to accomplish with the Interactive Storyteller, the digital tabletop board games in this section also combine the dynamics and intelligence of

¹Spiel Tisch AnReicherungs System (German for “game table augmentation system”)

¹<http://www.etc.cmu.edu/projects/surfacescapes/index.html>




Figure 2.10: SurfaceScapes Dungeons & Dragons (taken from <http://vimeo.com/7132858>)

computer games with the social advantages of traditional board games. By using a multi-touch table, like we have seen in these systems, we try to achieve interactive storytelling that resembles the social face-to-face setting of tabletop board games.

The view and interaction of all discussed tabletop board games are map-based. In the systems STARS MONOPOLY, KNIGHTMAGE and SURFACESCAPES (as well as FALSE PROPHETS [Mandryk & Maranan, 2002]), users can change the locations of characters by moving physical toys that represent the story characters across the map on the tabletop surface. To reinforce our resemblance with traditional board games, we decided to also investigate the use of tangible playing pieces that represent characters for physical interaction.

The SURFACESCAPES system and the STARS platform both use an additional upright screen to present additional images, text, video's or interface elements. For now, we decided not to add an upright screen to the Interactive Storyteller, to keep it as closely to traditional tabletop board games as possible and not to introduce another target for the users's attention besides the tabletop surface.

None of the tabletop board games discussed in this section are explicitly focussed on creating a story. That is the biggest difference that separates them from storytelling systems like ours and those discussed in section 2.1 and section 2.2.

Chapter 3

Design

This chapter introduces the design of the Interactive Storyteller. We start by looking at the requirements on which the design was based in section 3.1. In section 3.2 an overview of the functionality of the system and the interface design is discussed. In section 3.3 we discuss some of the findings from early user tests during the design process and in section 3.4 we look at the final changes that were made based on the pilot user test. In section 3.5 we look back at the requirements from section 3.1 and we discuss how they were implemented in the system. Unless otherwise noted, there are no differences between the touch-only and tangible+touch version of the system.

3.1 Requirements

Here we provide an overview of all requirements that we had for the design and implementation of the Interactive Storyteller. When a requirement was derived from literature, we provide a reference. When no reference is provided, we introduced the requirement ourselves. This section only serves as a summarised overview. It does not provide explanations about why particular requirements were adopted because then this section would become too extensive and cluttered to be used as a quick overview. Additional explanations can be found in the interface design text and in the reflection on the requirements at the end of this chapter. We divided the requirements into four categories: (1) requirements extracted from thesis introduction, (2) orientation and rotation requirements, (3) user experience requirements about immediate feedback and local agency, and (4) general multi-touch usability requirements.

Requirements derived from research goal and introduction

1. The interactive version of the Virtual Storyteller should provide a multi-user interface and a shared visual representation of the story world.
2. The system should be generic and domain-independent.
3. The platform should support face-to-face contact and social interaction.
4. Users should control characters in a sort of role-playing story game.
5. The system should be appropriate for children.

6. In line with traditional board games, the use of tangible playing pieces that represent characters for physical interaction should be investigated.

Orientation and rotation

7. People on all sides of the table should have a similar (equally good) view on the interface and the story world.
8. (Therefore) All interface elements should be recognisable and intelligible from all angles [Whalen, 2003].
9. Users should have the option to rotate items freely [Kruger et al., 2003].
10. When a user has positioned an item, that orientation must be maintained [Kruger et al., 2003].

Immediate feedback and local agency

11. Users should experience immediate feedback on the actions they perform [Mateas & Stern, 2005b].
12. Users should experience ‘local agency’ in the sense that they always have something to choose [Mateas & Stern, 2005b].

General multi-touch requirements

13. Global or shared interface elements should be presented in the centre of the table, so they are accessible for all users [Whalen, 2003].
14. Users should not be required to drag items over a long distance without the option to let them lift their fingers [Rick et al., 2009].
15. The tabletop should have a physical border around the table, so that users can lean without touching [Ryall et al., 2006].
16. Accidental touches should be taken into account. An accidental touch should not have a lasting effect [Ryall et al., 2006].
17. The interface should be suitable for different sizes of fingers [Ryall et al., 2006].

3.2 Interface design

The interface design¹ 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 [Schoning et al., 2008]. The Multi-Touch for Java (MT4j) framework² is used for multi-touch support.

¹An early version of the interface design was previously published in [Alofs et al., 2011].

²<http://www.mt4j.org/>



Figure 3.1: The tabletop used in the early user tests (see section 3.3).

3.2.1 Story world and view

The visual representation 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.

To match the Little Red Riding Hood domain and to appeal to our target users (see section 1.3), bright colours and cartoon-style images are used in the design (see Figure 3.5). Board games are however often played in a setting where children are joined by adults, or adults play with each other without children. Therefore a goal kept in mind was that the interface and interaction should also be appropriate for adults.

The story domain that we used is a variation of the original story of LRRH, with some additional locations and actions to increase choice. The story world 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). These locations are marked by blue circles on the map. 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. To increase choice, we added a few actions to the domain that are local in nature and have no effects on the story world (WatchBirds, EnjoySun, TakeNap, etc.). Characters can plan a series of 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).

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 position at the table of the user that currently has to choose an action can be determined, which is not the case with our hardware.

3.2.2 Moving characters

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 very much resembles the interaction offered by many familiar board games. 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 or drags a character image to an adjacent location, the blue circle of the



Little Red Riding Hood (Red)
...before she received her red cap



Grandma

Figure 3.2: Tangible playing pieces for the touch+tangible version.

destination location turns green to indicate that this is an allowed action. When the user moves the tangible or drags a character image 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 is not an allowed action.

Users might put tangibles or character images 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 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 or character image is not at that character's actual location in the story world, a dotted blue line is shown connecting the tangible or character image to its character's correct location. In the touch-only version, the system could automatically move character images back to their correct location. However, for comparison reasons this is not done. We wanted to keep the touch-only version as similar as possible to the touch+tangible version of the system to keep a 'fair' comparison.

3.2.3 Action Selection Interface

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 is usually done in computer games, but we consider it more important that the interface stays generic and is therefore not designed for any story world and set of actions in particular. Therefore, we decided to refrain from icon-based or other graphics-based ways for action selection, and use a flexible text-based approach.

Every time the turn goes to a new character, 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, and displays them in the ASI. The ASI can be seen in the centre of Figure 3.1 and in more detail in Figure 3.4. 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.

3.2.4 Story Areas

Users and spectators can read the results of actions that characters perform in the *story areas*, which are the two 'paper' *scrolls* that can be seen in the picture in Figure 3.1 and on the screenshot of the final interface design in Figure 3.5. If users want more or fewer lines of text to be visible in a story area, this can be achieved by touching the end of the scroll and rolling it up or down.

The sentences in the story areas are just straightforward representations of the actions executed so far. They read more like a sequential list of facts than as a nice coherent story because post-processing is not possible during story generation (but the narrator component of the Virtual Storyteller framework

can transform these facts into a fluent story afterwards [Theune et al., 2007]). These sentences are however very adequate for direct action feedback.

3.2.5 Resizing and positioning of interface elements

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 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 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 Shared Interfaces to Develop Effective Social skills (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 valuable screen estate.

3.3 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 (see Figure 3.1). 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 3.3). By fine-tuning some recognition parameters we managed to reduce these issues.

Another problem 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 often resulted in these users ending up confused and less immersed in the story. To address this issue, we decided to



Small fingertips were badly recognised. Hovering palms did get recognised.

Figure 3.3: Recognition problems in the setup used during the early user tests.

offer the same action feedback information in another (complementary) modality by vocalising it with Loquendo¹ text-to-speech (TTS), while keeping the story areas as a time-independent source of the same information about the story.

After introducing the new modality sound, we decided to also allow the option of adding action specific sounds. Although domain specific, associating actions with sounds is a very 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]. We decided not to add visual action feedback like additional images or animations to represent (the result of) actions. Creating these would require a lot of time, which was not available and we did not want to make the system more domain specific.

Also, we discovered that users often lost track of turn-taking. We solved this 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 3.4).

3.4 Pilot user test

A pilot user test involving two adult users was performed to ensure that the system was ready for our evaluation experiment (described in chapter 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. We got the opportunity to use the EVOLUCE ONE²³ multi-touch table owned by T-Xchange, a company located close to the university campus. This multi-touch table uses a *full HD* LCD screen instead of tweaking and calibrating a projector image reflected by two mirrors on an opaque surface (which the hardware in Figure 3.1 does). Instead of just one camera, it uses the merged image from three camera’s to cover the entire width of the tabletop surface, which is more accurate. Besides these advantages, it can also recognise smaller fiducial markers, so the base of the tangibles can be smaller (approximately one square inch). The use of this multi-touch table resulted in a severe reduction of the number of technical issues and recognition problems. The pilot user test did reveal a few remaining

¹<http://www.loquendo.com/en/demo-center/interactive-tts-demo/>

²http://www.evolute.com/en/hardware/multi-touch_table.php

³<http://www.evolute.com/downloads/docs/Multi-touch-Table-English-2012.pdf>

problems with the interface, which were small and could easily be fixed.

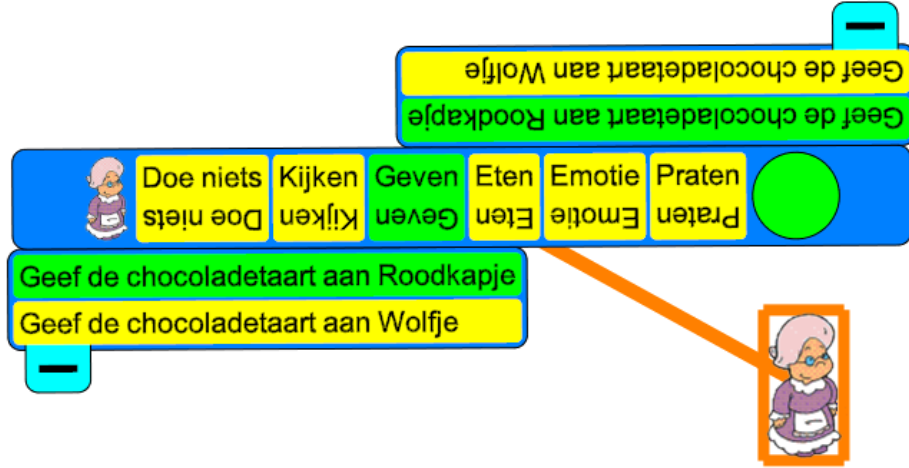


Figure 3.4: ASI connected to character that has the turn by a solid line and character image in ASI.

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 3.4). 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 (see Figure 3.5).

Finally, we changed the ASI so that it displayed a message like “CharacterX is trying to make a plan...” whenever a system-controlled character was busy making a plan for longer than one second. System-controlled characters use a *partial-order planner* to reach their goals [Russell & Norvig, 1995]. 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 the maximum search depth is reached, because there is no possible plan to satisfy Wolfs hunger. Reaching this maximum search depth takes 18 seconds on average, but occasionally even double that time. Unfortunately we could not find a way to abort the partial-order planner that runs inside the *Prolog* knowledge base of the Virtual Storyteller framework after a specified amount of 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. For now, adding a message to the ASI to inform the user about what they are waiting for was the only feasible resolution.

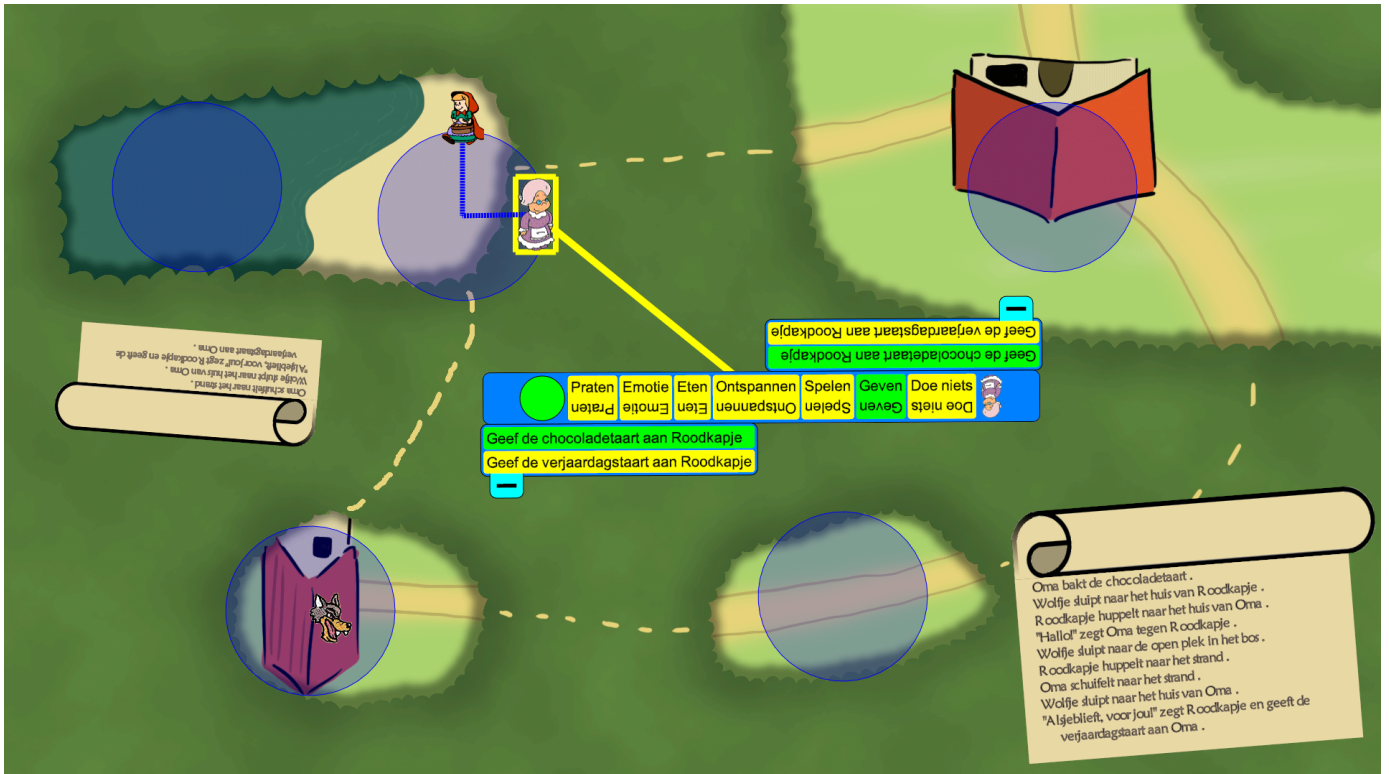


Figure 3.5: Screenshot of the final interface design (touch-only version).

3.5 Reflection on requirements

In this section we look back at the requirements from section 3.1 to see whether they are met in the system and how. The corresponding requirement numbers are displayed between brackets.

3.5.1 Requirements derived from research goal and introduction

The Interactive Storyteller uses a multi-user tabletop interface to display a shared visual representation of the story world (1). The system is generic and domain-independent because the only essential additions to enable an existing specification of a story domain to be used in the Interactive Storyteller are a file with coordinates and images for map and characters (2). The tabletop platform with board game metaphor was chosen because it is expected to support face-to-face contact and social interaction [Magerkurth et al., 2004] (3). Users interact with the story like they are actors that play the role of one of the characters, but without a director or script to tell them what to do (4). The early user tests helped to ensure that the system is appropriate for children (5). To investigate the use of tangible playing pieces to represent characters, two versions of the system were developed (6).

3.5.2 Orientation and rotation

Because the locations (houses) on the map are orientated one way and characters are orientated the other way, people on all sides of the table more or less have an equally good view on system (7). By not using orientation specific icons and by using two-sided text, all interface elements are recognisable and intelligible from all sides of the multi-touch table (8). Users can freely rotate the interface elements (ASI and story areas), but locations and characters currently have a fixed orientation (9). When the users have positioned the ASI or a story area, their location, size and orientation is maintained (10).

3.5.3 Immediate feedback and local agency

Users experience immediate visible and audible feedback on their actions by the text in the story areas, the text-to-speech engine and some sound effects (11). Eleven location-dependent actions were added to the LRRH domain to increase the number of available actions and ‘local agency’ (12).

3.5.4 General multi-touch requirements

At the start of each interaction session, the ASI is positioned in the centre of the tabletop (13). While users are dragging characters around, nothing changes when this dragging is momentarily interrupted by lifting their fingers for any amount of time (14). Both the tabletop from the early user tests as well as the tabletop used for the pilot and the final user tests had a physical border, so that users can lean without touching (15). Accidental touches have no irreversible effects because after touching an action, the round confirmation button also has to be touched before the selected action is executed (16). Users can scale the ASI and story areas, so the interface can be resized to suit any finger size (17).

Chapter 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 to what extent the system supports social interaction, we carried out a small-scale, qualitative user evaluation. The IRIS consortium performed very thorough research on the empirical assessment of the user experience in interactive storytelling [Klimmt et al., 2010]. However, the proposed evaluation method is a questionnaire, which is not very suitable for young children [Read & MacFarlane, 2006], which are our primary target user group (see section 1.3). We decided to mainly rely on observations using an observation scheme, as will be further discussed in section 4.2.

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 investigated the use of tangible playing pieces that represent the characters in a story. As discussed in subsection 3.2.2, 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 in some form or another 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.

We start by describing the procedure and experiment setup in section 4.1. After that the annotation method is discussed in section 4.2. In section 4.3 we look at the inter-annotator agreement. The results of the annotation and observations are discussed in section 4.4. In section 4.5 the answers to the interview questions are summarised. We conclude with an example of a created story in section 4.6.

4.1 Experiment setup and 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 4.1 provides an overview of the



Figure 4.1: Experiment setup and interaction (tangible version).

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 4.1: The four pairs of participants in our user evaluation.

The procedure was as follows. First, the *facilitator* (the researcher conducting the experiment) 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 particular direction. They were not specifically asked to talk or think aloud, because we were interested in their spontaneous behaviour while interacting with the system.

Due to our generic design, 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 users¹. The character controlled by AI should always be the same in order to eliminate any possible

¹When only one character would be controlled by the users, the resemblance with traditional board games would decrease, because in those games every user often controls his/her own playing piece. When all characters in the story world would be controlled by the users, the essential (intelligent) part of the storytelling framework is bypassed and just a passive simulation environment remains.

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

The child located nearest to Red’s house got assigned the character of Red to control; the child nearest to Grandma’s house 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 multi-touch table.

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.

The entire experiment 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 4.1. The recordings enabled easy observation of the children’s interaction with the system, what they said, the state of the story world, facial expressions, social behaviour, etc.

In principle, each interaction session lasted 15 minutes. Although the children continued playing without encouragement, 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).

4.1.1 Interview questions

Besides observation, part of the evaluation was an informal interview. The participants were interviewed together at the same time, which was recorded on video. There was not one single thing we wanted to learn from this interview. The purpose of several questions was only exploratory and some were merely a long shot. Not all questions were always answered, but most of the time follow-up questions were asked. Questions were reformulated when the children did not seem to understand them correctly.

After the children interacted with the first version of the system, they were asked three short questions:

1. What do you like about the system you just used?
2. What do you think is bad about the system you just used?
3. What do you think about the way you can move characters around?

The first two questions were mainly to break the ice and to give them the opportunity to freely speak their minds immediately after the first interaction session. After the second interaction session (with the other version of the system), they received some more questions. The first three questions are almost the same as after using the first version:

1. What do you like about the system you just used?
2. What do you think is bad about the system you just used?
3. (a) What do you think about the way you can move characters around?

- (b) In which version you think the way of moving characters was better?

Four questions focused on social behaviour:

4. In which version did the two of you collaborate the most?
5. In which version did the two of you talk the most?
6. In which version did the two of you have the most eye contact?
7. In which version did the two of you laugh the most?

The children were also asked which version was more fun and which version was easiest to operate:

8. Which version did you find the most fun?
9. Which version did you find the easiest?

We also asked about their attitudes towards group play and whether they would like to play again with one of the two versions:

10. Did you like playing this game with two people?
11. Would you also like to play this game alone?
12. Would you also like to play this game with more people? How many?
13. Would you like to play this game more often? Which version?

Finally, we asked if they had any experience with touch devices, which might have influenced their ease of interaction with, and expectations of, the multi-touch interface:

14. Have you ever used an iPad or other tablet? How often?
15. Have you ever used a smartphone? How often?

For the original Dutch interview questions see Appendix A.

4.2 Annotation method

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 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)¹, 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.

For our research we added an extra category to the SPC. The original SPC distinguishes four classes of social play (or ‘play domains’) with increasing levels of reciprocity and momentum in language as well as actions. In the original SPC

¹<http://www.rubin-lab.umd.edu/Coding%20Schemes/POS%20Coding%20Scheme%202001.pdf>

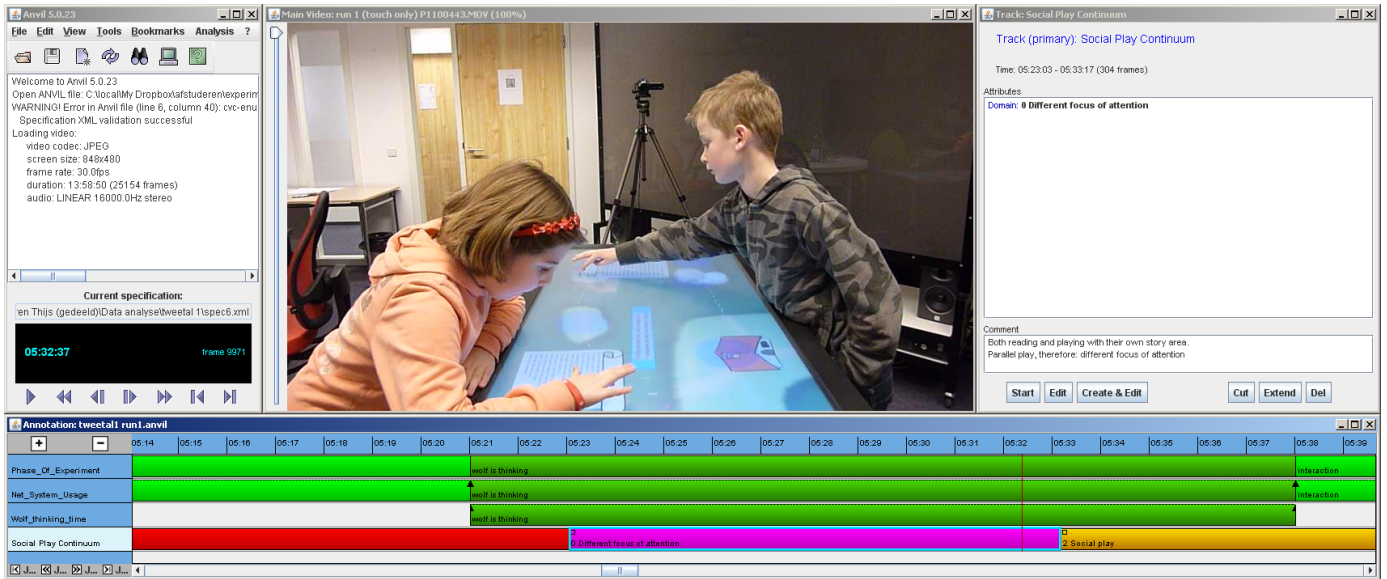


Figure 4.2: Annotation using the ANVIL 5 video annotation tool.

parallel play 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 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 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.

Because the time involved in fully annotating all the video recordings of the user test sessions, the observation scheme was not made more detailed than this. When considered necessary, more depth or detail could always be added afterwards.

In addition to the summary list above we used three SPC tables¹² by Pat Broadhead. These three tables describe the SPC domains on three different levels, from global to detailed. The table that describes detailed behaviours is very useful during annotation as it can be used to map observed actions and language to the corresponding SPC domains. Annotation was performed using the ANVIL³ [Kipp, 2001] video annotation tool (see Figure 4.2).

Annotators were required to put every video frame, where the system is running smoothly and no adult (facilitator) is talking, in a segment. Each segment should be classified in one of the 5 categories of the extended SPC (see list above). It is not possible for annotators to only annotate the clear and/or obvious moments, they should annotate all data.

We decided to use state coding instead of event coding. We did consider event coding (frequency-based approach), but state coding (duration-based approach) seemed more appropriate because of the small number of test sessions. We feel that without the possibility of good quantitative comparison “*cooperated x times*” is less informative than “*cooperated x% of the time*”. Also, it can be hard to determine whether for instance three consecutive moments of cooperation should be annotated as one, two or three ‘cooperations’.

Segmentation is free and up to each annotator. There is no minimum or maximum duration for a category time block and the annotator is free to determine the start and end time of each segment. This approach is more work for the annotators, but free segmentation offers more precision and it prevents annoying problems such as deciding on which is the most prominent activity/category within a predefined time block. Each segment can belong to only one category, so it can not have multiple category labels simultaneously. Two segments of the same category directly after each other, is considered the same as one large segment.

Participants are not annotated separately, because social behaviour is a thing between people and not something to be attributed to an individual. So, only one annotation track was used for each pair. With groups larger than two, another annotation approach might be more appropriate, but this depends on the research goals.

4.2.1 Additional annotation instructions

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. The additional annotation instructions are:

1. When several indications for multiple different categories are present, follow the strongest or most evident indication.
2. When several equally strong indications for multiple different categories are present, pick the ‘highest’(most social) corresponding category.
3. When in doubt whether something is a dialogue, **it is**.
4. When in doubt whether something is a “behavioural cluster”, **it is**.
5. When in doubt whether something is a smile, it is **not**.

¹<http://cw.routledge.com/textbooks/0415303397/resources/pdf/side1and2.pdf>

²<http://cw.routledge.com/textbooks/0415303397/resources/pdf/4domains.pdf>

³<http://www.anvil-software.de/>

6. No distinctions are made for the direction or target of a smile or laughter, because this is often too hard to discriminate.

4.3 Agreement analysis

One annotator used the revised annotation instructions to annotate all interaction sessions. One of the interaction sessions videos (different from the one used in the refinement of the annotation instructions) was annotated by a second annotator, to establish inter-annotator agreement.

		2 nd Annotator				
1 st Annotator	Different focus of attention	Different focus of attention	Associative play	Social play	Highly social play	Cooperative play
	Different focus of attention	4477	<i>3576</i>	120	0	0
	Associative play	724	39449	706	0	0
	Social play	937	<i>10352</i>	3984	0	0
	Highly social play	887	<i>9038</i>	<i>3728</i>	2754	0
	Cooperative play	0	473	1543	150	0

Table 4.2: **Confusion matrix.** Each time slice corresponds to 0,01 second in the annotated video. The length of the video was 13:49, which is 82898 time slices in total. So in example, a particular disagreement of 706 slices sums up to 7 seconds in total. The largest disagreements are marked in *italics*. Agreement is indicated by the bold numbers in the diagonal.

For determining inter-annotator agreement we used Brennan and Prediger’s free marginal kappa (K_M) (“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, 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 in our study, for comparison reasons we also provide Cohen’s kappa, which is 0.3389.

Annotating the level of social interaction is very subjective, which is also indicated by the disagreement in confusion matrix in Table 4.2. 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. The first annotated video was segmented into 85 fragments by the first annotator and into 46 segments by the second annotator. 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.

Presumably, a higher agreement could be reached by further refining the annotation instructions and more extensive training of the annotators, but given the small scale of the evaluation experiment we decided not to invest time in this.

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 4.4. For completeness, we provide the quantitative data that resulted from our annotation efforts in Appendix B.

4.4 Results of annotation and observations

For the qualitative analysis of the children’s interaction sessions, we used the annotated video segments that had been labelled by the first annotator as belonging to either the most social domains of the SPC (highly social play and cooperative play) or the least social category (different focus of attention). To enlarge the data set, we also looked at all the 32 fragments marked as those categories in the two videos annotated by the second annotator. In Table 4.3 the number of video fragments belonging to the most social and the least social domains are listed, as well as the total number of segments that were created during annotation.

Session	Version	HSBC	DFA	Total
Pair 1 run 1	touch only	9	11	102
Pair 1 run 2	tangibles	12	13	80
Pair 2 run 1	tangibles	5 (6)	20 (7)	85 (46)
Pair 2 run 2	touch only	7	11	85
Pair 3 run 1	touch only	11 (6)	15 (13)	98 (92)
Pair 3 run 2	tangibles	6	10	86
Pair 4 run 1	tangibles	14	5	70
Pair 4 run 2	touch only	18	6	88
Sum (all sessions)		82 (12)	91 (20)	694 (138)

Table 4.3: Number of video fragments belonging to Highly social behaviour and Cooperation (HSBC), Different focus of attention (DFA) and the total number of annotated video fragments. The results of the annotations by the second annotator are presented between brackets.

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 aspects to strengthen and which to avoid in order to improve the social aspects

of our system. After trying to determine all the reasons, we grouped them in clusters. Both assessing the reasons as well as grouping them into clusters was a fairly subjective process. We gained some interesting insights, which are presented in the following two subsections.

Because the results in the next two subsections come from a subjective process, we present observations, interpretations, thoughts and opinions all together. By doing this we keep things that belong together in one place, instead of pretending they are strictly separated by discussing them somewhere else.

The frequencies provided in the titles should only be regarded as an indication. These frequencies sum up to more than the number of video fragments (as presented in Table 4.3) because sometimes the behaviour in a particular video fragment has more than one reason. The order in which the reasons are presented is not based on the number of occurrences, but mainly on importance and interest.

No differences between the touch-only and touch+tangible version of the Interactive Storyteller could be found based on our observations and analysis of the video fragments. Therefore, we discuss everything together for both versions below.

4.4.1 Highly social behaviour and cooperation

The two highest levels of social behaviour from the SPC are highly social behaviour and cooperation. Inspection of the video segments where the children displayed behaviour in these two highly social categories revealed the following reasons or ‘social triggers’ in our setting. First, intended reasons for social behaviour will be discussed and after that we discuss more unintended reasons that however lead to highly social behaviour or cooperation.

Fun (13x)

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.

Interaction between the players’ characters (11x)

Highly social behaviour between players often originated from social interaction between their characters. Examples of this are the player characters giving each other cakes, talking to each other in the story world, or becoming angry at each other. 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 (10x)

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.

Acquiring help to obtain story goals (9x)

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.

Explaining character actions (10x)

After choosing an action for their character, players often explained the reason why they chose that particular action to the other player. Also, when they gained new insights about the story world or Wolf's behaviour, they 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 (20x) and/or reflections on Wolf's behaviour (6x)

For reasons explained in the last part of section 3.4, players often had to wait while Wolf was busy making a plan. On average, this took up around 27% of the total system interaction time of each test run. 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.

Providing help (18x)

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 for one player but not the other, or that did not work as expected.

Turn-taking (6x)

Players occasionally reminded the other that it was their turn. On the surface, this is cooperative behaviour, but in some cases it probably had a less social

underlying motivation, i.e., the player wanting to reduce the time spent waiting for the other to perform his/her turn.

Bugs, flaws and quirks of the system (7x)

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 ("Huh!? Red skips to Red's house? Why not *her* house?"). 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.

4.4.2 Different focus of attention

Besides the two highest levels of social behaviour from the SPC, we are also interested in the least social behaviour players display. 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. Having a different focus of attention is considered undesirable behaviour because a shared focus of attention is the first step towards social, highly social, and cooperative behaviour. We observed a few clear reasons for players having a different focus of attention in our setting.

Waiting for Wolf (36x)

Although waiting for Wolf often served as a trigger for (highly) social interaction (see subsection 4.4.1), it more frequently induced the participants to lose their shared focus of attention. When having to wait for Wolf, players often started looking around, started playing with different interface elements or started reading the story so far in their own story area. It seems that losing their shared focus of attention can be attributed to the (long) thinking times of Wolf.

Waiting for the other player (5x)

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 his/her 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 (24x)

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, kind of like staring while thinking about something, possibly the story or the system.

Playing around with the story area (32x)

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 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.

4.5 Interview results

The informal interviews (see subsection 4.1.1) that we held with the children were for a part aimed at finding out their preference for one or the other version of the system, their perception of social interaction when playing with either version, and their thoughts about the multi-player nature of the Interactive Storyteller.

Interview questions 1 and 2 resulted in many wild ideas and suggestions, but mainly about adding new actions or possibilities to the LRRH story world. However, because domain specific improvements are outside of the scope of the current research, we did not look into this any further.

In the remainder of this section we summarise the results from the other interview questions. Because the interviews were very informal and the children did not always answer each question, we found little use in presenting an overview of the ‘raw’ results. Instead, we summarise the outcome of the interview questions and discuss our main findings below. The results of the questions are grouped and reordered for coherence, but the question numbers have been added for easy reference.

4.5.1 Previous experience

[Question 14 and 15] We start with the last two questions, which were about the children’s 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 might have influenced their expectations and ease of interaction with the interface, but we did not investigate this any further.

4.5.2 Version preference

[Question 3, 8, 9 and 13] When asked about their preference for either of the two versions of the system, the children’s preference for a 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 found it annoying 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 does not matter which version of the system players use, as long as the movement of characters is working properly and players do not experience any trouble with it.

4.5.3 Social interaction

[**Question 4, 5, 6, 7 and 8**] 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. 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, follow-up questions revealed that this was practically 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, follow-up questions revealed that this 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, which in turn provided them with more grounds for social interactions such as communication and collaboration.

4.5.4 Desire to play again

[**Question 13**] 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 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.

4.5.5 Desire to play alone or together

[**Question 10, 11, 12**] Most children answered that they rather would like to play with the system together instead of 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.

4.6 Example of created story

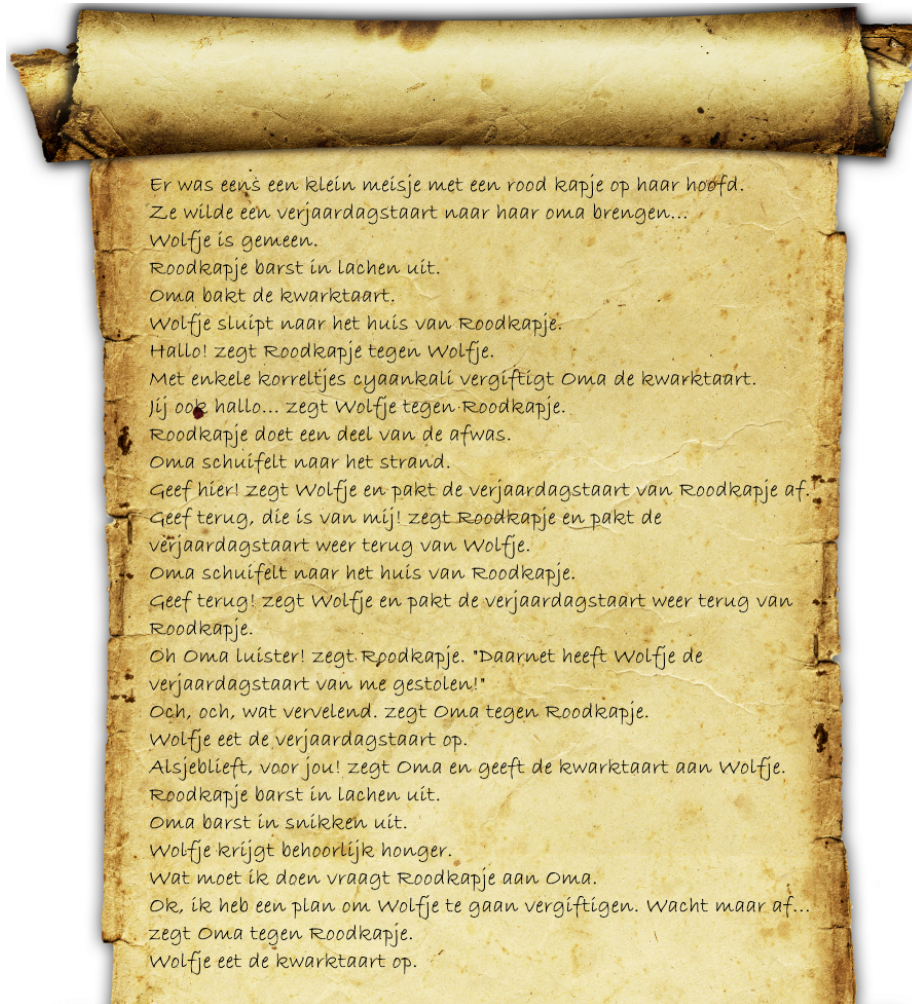


Figure 4.3: Example of a created story by the Dutch children in the experiments.

To illustrate the kind of stories that were created in the final user tests, this section provides a part of one of the stories created by the children. The sentences originate from the events as they happened, so without post-processing the story afterwards, which is optional in the Virtual Storyteller framework with the narrator component to increase quality and readability of the story [Theune et al., 2007].

The actions chosen by the children sometimes may seem somewhat random (for instance, in the example below, Grandma bursting out in tears without any reason), but sometimes they were based on something that happened outside of the system, or they were just chosen for the sake of fun of acting strange or unexpectedly.

The Interactive Storyteller also uses a generic approach for languages. To

get the English equivalent of the children's story in Figure 4.3, it was reenacted step by step by performing the same sequence of actions in the English version of our system. The story of the first test run of the fourth pair was like this:

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

Chapter 5

Discussion

In this chapter we discuss some of our findings and the ideas and thoughts we have about them in a bit more detail. We will discuss things related to turn-taking in section 5.1. We will discuss our findings on the use of tangibles in section 5.2. We will discuss things related to cooperative behaviour and goals in section 5.3.

5.1 Social interaction imposed by turn-taking

The long period it often took for Wolf to realise there was no plan to satisfy the adopted goal, turned out to have unexpected benefits. As expected, these long waiting times resulted in the players getting distracted while waiting, but in many cases they served as a trigger for social interaction in the form of story-related discussions. This was helped by the fact that while waiting for Wolf, both players had no other occupation. Our findings suggest 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 those periods none of the players can perform any actions and they are implicitly (or maybe even explicitly) invited to discuss past and future story developments.

Games that use turn-taking inherently introduce waiting time for each of the participating players. In our experiments, having to wait for the other player to finish his or her turn 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 player did, simply because they could not do much else at that time. This (presumably) led to higher involvement in the actions the other player chose, and thus provided more options for social interaction.

Interesting research on turn-taking 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]. We are very interested in the upcoming results of this research, but at least until then, we stick to our strict turn-taking approach.

5.2 Tangibles

We did not find an influence of the interaction style (tangibles or touch only) on the players' social behaviour. Not in our observations of the children's interaction sessions and not in the results of the informal interviews we carried out. We did also not find any other advantages for tangible playing pieces to represent player-controlled characters in our storytelling board game. The fact that we did not find any influence of the interaction style in our experiments does however not mean that we can conclude that there *is* no influence. Nor can we conclude anything about the interaction styles in general.

In order to be able to make a fair comparison between tangibles+touch and touch-only in our experiment, we minimised the differences between both versions. Ideas for functionality that were only feasible to implement in one of the two versions were not used during the user experiments. By using this highest common denominator approach we prevented that any differences in the test results would originate from inequalities in functionality.

Another approach would have been to maximise the differences and to take full advantage of all the different properties and benefits of tangibles and touch-only images. In that case it would however never be possible to draw the conclusion that any (social) outcome is the result of the use (or non-use) of tangibles. Any particular outcome could always also be attributed to one or several of the differences in functionality. We were not interested in any specific functionality set, but wanted to find out whether the use of tangibles had some clear added value in itself for our system.

5.3 Cooperation and goals

Cooperative behaviour mostly occurred when the players helped each other by explaining or demonstrating the interface. This is something that should no longer be necessary when the interface is better explained or further improved. Cooperation sometimes also took the form of working towards a particular goal together, and this 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, 2005b]. 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.

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. In relation to goal-orientated 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. Other ‘flow’ elements are immediate feedback on actions and a sense of control, which are both supplied by the Interactive Storyteller.

Chapter 6

Conclusions

We succeeded in our goal to design, implement and evaluate a multi-user domain-independent interface to be used with our interactive storytelling framework. Social interaction is supported by presenting the visual representation of the story world on a shared multi-touch table. We found that the turn-based nature of our framework increases the number of opportunities for social interaction. We also discovered that having to wait for the planning algorithm of a computer controlled character is not only a bad thing, because it also induces social interaction.

To reinforce the resemblance with traditional board games, we investigated the use of tangible playing pieces that represent the characters in the story. To answer the question whether the use of tangibles in our system has advantages over a touch-only approach, we observed the videos of the interaction sessions and tried to find differences in social behaviour. We also asked the users their preference after using both versions. We did not find clear differences in social behaviour based on our user experiments. According to the answers from the users it seems that it does not matter which version of the system they use, as long as the movement of characters is working properly and players do not experience any trouble with it.

The use of tangibles imposed quite a few restrictions on our system and constrained its flexibility. 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 anywhere on the map, the system should never rotate, zoom or tilt the map, because it can not physically move the tangibles. Because of these limitations and because we did not find clear advantages for the use of tangibles, we suggest that the next version of the Interactive Storyteller should focus on touch-only interaction.

The Social Play Continuum (SPC) offers a small number of interaction categories that cover the entire spectrum from non-social to highly social behaviour and thereby provides a simple and decent way to analyse social interaction. Although we did not succeed in achieving a high inter-annotator agreement (which is not uncommon in the analysis of social behaviour), we gained some interesting insights thanks to the SPC. From the results of the SPC annotation and our observations we determined the triggers for displaying the most social and the least social behaviour (section 4.4).

- Triggers for highly social behaviour and cooperation are: funny actions or sounds, interaction between the players' characters, planning for the next actions or goal, acquiring help to obtain story goals, explaining character actions, waiting for Wolf, reflections on Wolf's behaviour, providing help, turn-taking, and bugs or flaws or quirks of the system.
- Triggers for having a different focus of attention are: waiting for Wolf, waiting for the other player, playing around with the story area, or 'just' being distracted.

From the interview results (section 4.5) we learned that:

1. All children in our research had previous experience with multi-touch devices and often played games on smartphones or tablets at home or somewhere else.
2. When children had a preference for one of the two versions, this was always based on having trouble with the other version. Overall, there was no version that was preferred over the other.
3. There were no differences in experienced social behaviour for both versions according to the children. All eight children said they had laughed the most while using the second version, whichever version that was.
4. All children would like to have played with the system again afterwards.
5. Most children rather would like to play with the system together instead of alone, preferably with 3-5 people and in a family setting.

By using a multi-touch table we aimed to achieve interactive storytelling that resembles the social setting of traditional tabletop board games. Like existing digital tabletop board games, we tried to combine the dynamics and intelligence of computer games with the social advantages of traditional board games. To our knowledge, the Interactive Storyteller is the first AI-based interactive storytelling system that uses the social aspects of traditional tabletop board games.

Chapter 7

Recommendations

In this final chapter we provide some recommendations for future work. Besides the obvious suggestion of repeating the same experiment with a much larger sample size, we discuss some ideas for follow-up experiments in section 7.1. We give suggestions for customisation in section 7.2, to improve scalability in section 7.3 and about visual feedback in section 7.4. Our ideas for introducing moments of reflection are explained in section 7.5. Finally, we present the idea to use a tablet as a gameboard in section 7.6.

7.1 New experiments

For the user experiments as discussed in this thesis, we did not want to provide the users with particular instructions or a goal to achieve. We did not tell them to cooperate and we did not instruct them to create a nice, good, funny or coherent story. In fact, we did not instruct them to create any story at all, we only said that the system could be used to create a story. We did not want to steer them in any direction, because this would influence the outcome of the experiments. For a follow-up experiment it would be interesting to investigate what would happen if the users receive a particular goal. After a few minutes of playing around with the system and noticing that Wolf always tries to steal your cakes, we could instruct the users to start over and create a story in which they get back at Wolf. It would also be interesting to see how the users would behave on a social level and how the story would develop if the users receive explicit instructions to cooperate to create a (nice, good, funny or coherent) story together. Summarised, in a follow-up experiment we want to provide the users with clear goals (which is in line with the *flow* theory of Csikszentmihalyi [1991]) to improve the user experience and increase social interaction.

Another idea for follow-up experiments is to investigate the influence of group composition and group size. Because of our ‘board game in a family setting’ idea, we are interested in larger groups and more diversity in age. We would like to test the Interactive Storyteller with actual families to see how this influences the story development and the social interaction between users.

We are also interested whether users prefer to play one character each, like actors playing a role, or rather choose the actions for all characters together in deliberation. Each user can play one character (like in the already performed

experiments), but collaborative control is also possible, with users cooperating to make decisions for all the human-controlled characters together. This alternative approach could be researched by just looking what users would do spontaneously, or by providing explicit instructions to discuss and choose all actions together. These instructions could possibly be enforced by requiring a multi-user touch gesture on the tabletop to confirm each chosen action. Experiments with more forced cooperation are however also possible, for instance an experiment in which there is just one human controlled character that has to be controlled by two users.

7.2 Customisation

Because we have found no indications that there are advantages of tangibles to represent characters in our system, we decided that the next prototype of the Interactive Storyteller should not use tangible playing pieces, but should be based on touch-only user input. An advantage of using images to represent characters instead of tangibles is that it offers the possibility of easy 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 preferences of 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 available in many contemporary computer games is relatively easy and could provide in the need, expressed by some children, to make their character more “their own”.

7.3 Scalability

The scalability of the action selection interface (ASI) for larger domains with more choice should be further improved. When the character Red has four cakes and there are two characters at the same location as her to give the cakes to, this already results in a list with eight possible ‘give’-actions. Considering larger domains with even more objects and characters, this can lead to an explosion of possibilities. The possible arguments of every action predicate are already available in our storytelling framework. Currently the possible actions in the ASI are only grouped by action type. When there are many instantiations of a particular action (like in the ‘give’ example just mentioned), the number of options to choose from should be reduced by first requesting the user to select the *patient* of the action before displaying the list of (remaining) possible instantiations of the action. When there are still a lot of instantiations of a particular action left, the same approach could be used for the ‘recipient’, ‘instrument’ and possibly even other *thematic roles* before displaying the list.

7.4 Visual feedback

The visual representation of the story world currently only changes after movement actions. A possible direction for future research is to add more visual feedback about the state of the story world and the results of selected actions. For example, by adding object inventories that contain small icons of the things that are carried by a character, or the use of animations as action feedback. We expect that this will improve the user experience and make the interaction more engaging, but because such additions are domain specific, it will be more work to enable a new or existing story domain to be used in the Interactive Storyteller. Allowing the possibility to add engaging visual elements while keeping the framework generic remains an open challenge.

7.5 Moments of reflection

As mentioned in the discussion, even when more efficient planning for system-controlled characters is implemented, it might still be a good idea to build in some forced moments of reflection. Even when Wolf already planned his next action, we can still pretend he continues thinking. This additional waiting time provides the users with a moment to discuss past and future story developments while they both have nothing else to do. This additional waiting time can be random or have a fixed duration. An even better idea would be to make it dynamic, based on computer vision or audio analysis of the users. After the minimal amount of thinking time, the system tests whether the users are looking at each other, have a shared focus of attention, are talking, or whether they are silent and have a different focus of attention. Based on this audio and/or video analysis, the ‘thinking’ time of the computer-controlled character could be dynamically prolonged to sustain any possible ongoing social interaction. Tone and pitch analysis could be used to determine the presence of frustration, which could be an indication to stop the ‘thinking’ right away and alert the users that Wolf has made a plan. This thinking time could also actively be used by the system to provide assignments or explanations, for instance in a serious game that trains and educates autistic children about social skills.

7.6 Tablet board game

Finally, in the future we would like to bring the Interactive Storyteller experience to the new generation of tablet computers. 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 gameboard of a traditional board game. Using a tablet as a gameboard in the living room would bring us closer to our vision of the the Interactive Storyteller as a ‘family board game with benefits’.

Glossary

- agent** an autonomous software entity which observes and acts upon an environment and directs its activity towards achieving goals [Russell & Norvig, 1995]. 1, 2, 6–9, 29, 51
- AI-based** based on AI. i, 1, 2, 5–10, 13, 29, 47
- antagonist** Greek for: opponent, competitor, enemy, rival. The character that opposes the main characters. The character against which the protagonist must contend. (see: <http://en.wikipedia.org/wiki/Antagonist>). 30
- branched narrative** see: scripted narrative. 10
- emergent narrative** is narrative generated by running a simulation in which character agents interact. i, 1, 7, 8, 10
- facilitator** the researcher conducting the experiment. 29, 30, 33, 38
- fiducial** a visual printed pattern that is recognisable by computer vision software and used to identify and track tangibles. 19, 24
- finite state machine** a mathematical model used to design computer programs. It can be viewed as a function which maps an ordered sequence of input events into a corresponding sequence of output events (<http://definitions.dictionary.net/finite%20state%20machine>). 6
- flow** a state of optimal enjoyment, or: “the state in which people are so involved in an activity that nothing else seems to matter; the experience itself is so enjoyable that people will do it even at great cost, for the sheer sake of doing it.” [Csikszentmihalyi, 1991]. 44, 45, 48
- full HD** the maximum high-definition resolution (1920 x 1080 pixels). 24
- heuristic search planner** a planner that searches state space with a heuristic that can be extracted automatically from the problem encoding, and then combined with standard search algorithms (source: <http://scom.hud.ac.uk/planet/repository/heuristic.html>). 7
- knowledge base** a database for knowledge management and reasoning about facts. 22, 25
- multi-touch** something using or capable of recognising more than one touch point at the same time. i, ii, 2, 3, 11, 12, 16–19, 21, 23, 24, 27, 30, 31, 39, 40, 46, 47

- partial-order planner** an algorithm or program which will construct a plan and searches for a solution. The input is the problem description, consisting of descriptions of the initial state, the goal and possible actions (source: http://en.wikipedia.org/wiki/Partial-order_planning). 25
- patient** “the semantic role of an entity that is not the agent but is directly involved in or affected by the happening denoted by the verb in the clause” (source: <http://wordnetweb.princeton.edu/perl/webwn?s=patient>). 49
- Prolog** “a general purpose logic programming language associated with artificial intelligence and computational linguistics” (source: <http://en.wikipedia.org/wiki/Prolog>). 25
- scripted narrative** narrative in which all choice options and variations of possible stories are explicitly written out in advance. 7, 51
- scroll** the roll-up paper metaphor used for the story areas. 22
- serious game** a game designed for a primary purpose other than pure entertainment (see: http://en.wikipedia.org/wiki/Serious_game). 8, 10, 13, 50
- story area** the two pieces of ‘paper’ that display the story generated so far. They can be seen on the screenshot in Figure 3.5. 22–24, 27, 38, 39, 45, 47
- tangible** a touchable physical object to interact with a system. i, 3, 9, 11–13, 18, 19, 21, 22, 24–26, 28, 36, 40, 43, 44, 46, 49, 51, 63
- thematic role** “the role that a noun phrase plays with respect to the action or state described by a sentence’s verb” (source: http://en.wikipedia.org/wiki/Thematic_relation). 49

Acronyms

- AI** artificial intelligence. 9, 14, 29, 30, 51
- ASI** action selection interface. 22–25, 27, 49
- DEIRA** Dynamic Engaging Intelligent Reporter Agent. 1
- DFA** Different focus of attention. 35
- FearNot!** Fun with Empathic Agents Reaching Novel Outcomes in Teaching.
6, 7, 10
- HMI** Human Media Interaction. ii, 1
- HSBC** Highly social behaviour and Cooperation. 35
- IRIS** Integrating Research in Interactive Storytelling. 28
- LCD** Liquid Crystal Display. 24
- LRRH** Little Red Riding Hood. 3, 10, 20, 21, 25, 27, 39, 44
- MT4j** Multi-Touch for Java. 19
- NLE** narrative-centred learning environment. 6–8, 10, 13
- ORIENT** Overcoming Refugee Integration with Empathic Novel Technology.
8–10
- PLUNDER** Pirates Looming in Unscripted Narrative: towards Dramatic
Emergent Role-play. 3
- POS** Play Observation Scale. 31
- RFID** Radio Frequency IDentification. 9
- RPG** role-playing game. 8
- SIDES** Shared Interfaces to Develop Effective Social skills. 23
- SPC** Social Play Continuum. 31, 33, 35–38, 46
- STARS** Spiel Tisch AnReicherungs System (German for “game table augmen-
tation system”). 15, 17
- TTS** text-to-speech. 24

Bibliography

- Alofs, T., Theune, M., & Swartjes, I. (2011). A tabletop board game interface for multi-user interaction with a storytelling system. In *Proceedings of the 4th International Conference on Intelligent Technologies for Interactive Entertainment (INTETAIN 2011)*, volume 78 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering (LNICST)*, (pp. 123–128). Springer.
[Referenced in this thesis from page: 19]
- Alves, A., Lopes, R., Matos, P., Velho, L., & Silva, D. (2010). Reactoon: Storytelling in a tangible environment. In *Proceedings of the 3rd IEEE International Conference on Digital Game and Intelligent Toy Enhanced Learning (DIGITEL 2010)*, (pp. 161–165).
[Referenced in this thesis from page: 12, 13]
- Aylett, R. (1999). Narrative in virtual environments - towards emergent narrative. In *Proceedings of the 16th National Conference on Artificial Intelligence (AAAI 1999)*, Working notes of the Narrative Intelligence Symposium, (pp. 83–86). AAAI.
[Referenced in this thesis from page: 1]
- Aylett, R., Louchart, S., Dias, J., Paiva, A., & Vala, M. (2005). FearNot! - an experiment in emergent narrative. In *Proceedings of the 5th International Working Conference on Intelligent Virtual Agents (IVA 2005)*, volume 3661 of *Springer LNCS*, (pp. 305–316).
[Referenced in this thesis from page: 6]
- Aylett, R., Vala, M., Sequeira, P., & Paiva, A. (2007). FearNot! - an emergent narrative approach to virtual dramas for anti-bullying education. In *Proceedings of the 4th International Conference on Virtual Storytelling (ICVS 2007)*, volume 4871 of *Springer LNCS*, (pp. 202–205).
[Referenced in this thesis from page: 6]
- Broadhead, P. (2004). *Early years play and learning: developing social skills and cooperation*. Routledge.
[Referenced in this thesis from page: 31]
- 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.
[Referenced in this thesis from page: 31]

Bibliography

- 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.
[Referenced in this thesis from page: 3, 31]
- Cappelletti, A., Gelmini, G., Pianesi, F., Rossi, F., & Zancanaro, M. (2004). Enforcing cooperative storytelling: First studies. In *Proceedings of the 4th IEEE International Conference on Advanced Learning Technologies (ICALT 2004)*, (pp. 281–285).
[Referenced in this thesis from page: 11, 13]
- Charles, F., Pizzi, D., Cavazza, M., Vogt, T., & André, E. (2009). Emoemma: emotional speech input for interactive storytelling. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, (pp. 1381–1382). International Foundation for Autonomous Agents and Multiagent Systems.
[Referenced in this thesis from page: 7]
- 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.
[Referenced in this thesis from page: 2]
- Coldewey, D. (2010). Hands-on: D&D on the Microsoft Surface. <http://techcrunch.com/2010/02/10/hands-on-dd-on-the-microsoft-surface/>.
[Referenced in this thesis from page: 16]
- Csikszentmihalyi, M. (1991). *Flow: The psychology of optimal experience*. Harper Perennial.
[Referenced in this thesis from page: 2, 44, 48, 51]
- 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 (UIST 2001)*, (pp. 219–226). ACM.
[Referenced in this thesis from page: 16]
- Gal, E., Bauminger, N., Goren-Bar, D., Pianesi, F., Stock, O., Zancanaro, M., & Weiss, P. (2009). Enhancing social communication of children with high-functioning autism through a co-located interface. *Ai & Society*, 24(1), 75–84.
[Referenced in this thesis from page: 11]
- Helmes, J., Cao, X., Lindley, S., & Sellen, A. (2009). Developing the story: Designing an interactive storytelling application. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS 2009)*, (pp. 49–52).
[Referenced in this thesis from page: 11, 12, 13]
- Hornecker, E. & Buur, J. (2006). Getting a grip on tangible interaction: a framework on physical space and social interaction. In *Proceedings of the SIGCHI conference on Human Factors in computing systems: Designing for Tangible Interactions (CHI 2006)*, (pp. 437–446). ACM.
[Referenced in this thesis from page: 3]

Bibliography

- 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 2011)*, (pp. 199–206). IEEE.
[Referenced in this thesis from page: 2]
- Kipp, M. (2001). Anvil - a generic annotation tool for multimodal dialogue. In *Seventh European Conference on Speech Communication and Technology (Eurospeech 2001)*, (pp. 1367–1370).
[Referenced in this thesis from page: 33]
- Klimmt, C., Roth, C., Vermeulen, I., & Vorderer, P. (2010). The empirical assessment of the user experience in interactive storytelling: Construct validation of candidate evaluation measures. Technical report, IRIS Consortium (Integrating Research in Interactive Storytelling).
[Referenced in this thesis from page: 28]
- Knoppel, F., Tigelaar, A., Oude Bos, D., & Alofs, T. (2008). DEIRA: A Dynamic Engaging Intelligent Reporter Agent. In *Proceedings of the 20th Belgian-Netherlands Conference on Artificial Intelligence (BNAIC 2008)*, (pp. 393–394).
[Referenced in this thesis from page: 1]
- Korhonen, H., Montola, M., & Arrasvuori, J. (2009). Understanding playful user experience through digital games. In *Proceedings of the 4th international Conference on Designing Pleasurable Products and Interfaces (DPPI 2009)*, (pp. 274–285).
[Referenced in this thesis from page: 2, 49]
- 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 *Proceedings of the 3rd Workshop on Multimodal Output Generation (MOG 2010)*, (pp. 15–25).
[Referenced in this thesis from page: 8, 9]
- Kriegel, M., Lim, M., Nazir, A., Aylett, R., Cawsey, A., Enz, S., Rizzo, P., & Hall, L. (2008). ORIENT: An inter-cultural role-play game. In *Proceedings of the 5th International Conference on Narrative and Interactive Learning Environments (NILE 2008)*, (pp. 70–74).
[Referenced in this thesis from page: 8]
- 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 (Group 2003)*, (pp. 369–378). ACM.
[Referenced in this thesis from page: 19, 23]
- Lazzaro, N. (2004). Why we play games: Four keys to more emotion without story. <http://www.xeodesign.com/xeodesign-whyweplaygames.pdf>.
[Referenced in this thesis from page: 2]
- Lazzaro, N. (2008). The four fun keys. In K. Isbister & N. Schaffer (Eds.), *Game Usability: Advancing the Player Experience* chapter 20, (pp. 315–344).

Bibliography

- Morgan Kaufmann.
[Referenced in this thesis from page: 2]
- Machado, I., Prada, R., & Paiva, A. (2000). Bringing drama into a virtual stage. In *Proceedings of the 3rd International Conference on Collaborative Virtual Environments (CVE 2000)*, (pp. 111–117). ACM.
[Referenced in this thesis from page: 7]
- Macionis, J. J. & Gerber, L. M. (2010). *Sociology* (Seventh Canadian ed.). Pearson Canada.
[Referenced in this thesis from page: 2]
- Magerkurth, C., Engelke, T., & Memisoglu, M. (2004). Augmenting the virtual domain with physical and social elements: towards a paradigm shift in computer entertainment technology. *Computers in Entertainment (CIE)*, 2(4), 163–172.
[Referenced in this thesis from page: 2, 15, 26]
- 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).
[Referenced in this thesis from page: 15, 21]
- Magerkurth, C., Stenzel, R., & Prante, T. (2003). STARS - a ubiquitous computing platform for computer augmented tabletop games. In *Video Track and Adjunct Proceedings of the Fifth International Conference on Ubiquitous Computing (UBICOMP 2003)*.
[Referenced in this thesis from page: 15]
- Malone, T. & Lepper, M. (1987). Making learning fun: A taxonomy of intrinsic motivations for learning. *Aptitude learning and instruction*, 3(3), 223–253.
[Referenced in this thesis from page: 2, 24, 49]
- Mandryk, R. & Maranan, D. (2002). False prophets: exploring hybrid board/video games. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems: Changing our world, changing ourselves (CHI 2002)*, (pp. 640–641).
[Referenced in this thesis from page: 15, 17, 21]
- Mateas, M. & Stern, A. (2003). Façade: An experiment in building a fully-realized interactive drama. In *Game Developers Conference: Game Design Track (GDC 2003)*. Literature, Communication and Culture and College of Computing, Georgia Tech.
[Referenced in this thesis from page: 5]
- Mateas, M. & Stern, A. (2005a). Procedural authorship: A case-study of the interactive drama façade. *Digital Arts and Culture (DAC). Digital Experience: Design, Aesthetics, Practice*.
[Referenced in this thesis from page: 6]
- Mateas, M. & Stern, A. (2005b). Structuring content in the façade interactive drama architecture. In *Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2005)*, (pp. 93–98). AAAI

Bibliography

Press, Menlo Park, CA.

[Referenced in this thesis from page: 19, 44]

Mazalek, A. & Davenport, G. (2003). A tangible platform for documenting experiences and sharing multimedia stories. In *Proceedings of the 2003 ACM SIGMM workshop on Experiential Telepresence (ETP 2003)*, (pp. 105–109). ACM.

[Referenced in this thesis from page: 12, 13]

Mazalek, A., Reynolds, M., & Davenport, G. (2007). The TViews table for storytelling and gameplay. In C. Magerkurth & C. Röcker (Eds.), *Concepts and Technologies for Pervasive Games: A Reader for Pervasive Gaming Research*, volume 1 (pp. 1–29). Shaker Verlag.

[Referenced in this thesis from page: 12]

Pape, J. & Graham, T. (2010). Coordination policies for tabletop gaming. *Graphics Interface Poster*, 24–25.

[Referenced in this thesis from page: 43]

Piper, A., O’Brien, E., Morris, M., & Winograd, T. (2006). SIDES: a cooperative tabletop computer game for social skills development. In *Proceedings of the 20th anniversary conference on Computer supported cooperative work (CSCW 2006)*, (pp. 1–10).

[Referenced in this thesis from page: 23]

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

[Referenced in this thesis from page: 7]

Prada, R., Machado, I., & Paiva, A. (2000). Teatrix: Virtual environment for story creation. In *Proceedings of the 5th International Conference on Intelligent Tutoring Systems (ITS 2000)*, volume 1839 of *Springer LNCS*, (pp. 464–473).

[Referenced in this thesis from page: 7]

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*, volume 2363 of *Springer LNCS*, (pp. 913–922).

[Referenced in this thesis from page: 8]

Randolph, J. (2005). Free-marginal multirater kappa (multirater κ_{free}): an alternative to Fleiss fixed-marginal multirater kappa. In *Joensuu Learning and Instruction Symposium*, (pp. 1–20).

[Referenced in this thesis from page: 34]

Read, J. & MacFarlane, S. (2006). Using the fun toolkit and other survey methods to gather opinions in child computer interaction. In *Proceedings of the 2006 conference on Interaction design and children (IDC 2006)*, (pp. 81–88). ACM.

[Referenced in this thesis from page: 28]

Bibliography

- Rick, J., Rogers, Y., Haig, C., & Yuill, N. (2009). Learning by doing with shareable interfaces. *Children, Youth and Environments*, 19(1), 321–342.
[Referenced in this thesis from page: 19]
- Rubin, K. (1982). The play observation scale. *Child Development*, 53(3), 651–657.
[Referenced in this thesis from page: 31]
- Russell, S. & Norvig, P. (1995). *Artificial intelligence: A modern approach*. Prentice Hall (Englewood Cliffs, NJ).
[Referenced in this thesis from page: 25, 51]
- Ryall, K., Morris, M., Everitt, K., Forlines, C., & Shen, C. (2006). Experiences with and observations of direct-touch tabletops. In *Horizontal Interactive Human-Computer Systems, 2006. TableTop 2006. First IEEE International Workshop on*, (pp. 8–pp). IEEE.
[Referenced in this thesis from page: 19]
- Schoning, J., Brandl, P., Daiber, F., Echter, F., Hilliges, O., Hook, J., Lochtefeld, M., Motamedi, N., Muller, L., Olivier, P., et al. (2008). Multi-touch surfaces: A technical guide. Technical report, TUMI0833: Technical Reports of the Technical University of Munich.
[Referenced in this thesis from page: 19]
- Swartjes, I. & Theune, M. (2008). The Virtual Storyteller: Story generation by simulation. In *Proceedings of the 20th Belgian-Netherlands Conference on Artificial Intelligence (BNAIC 2008)*, (pp. 257–265).
[Referenced in this thesis from page: 1, 3]
- Swartjes, I. & Theune, M. (2009). Iterative authoring using story generation feedback: Debugging or co-creation? In *Proceedings of the 2nd International Conference on Interactive Digital Storytelling (ICIDS 2009)*, volume 5915 of *Springer LNCS*, (pp. 62–73).
[Referenced in this thesis from page: 3, 44]
- Theune, M., Faas, S., Nijholt, A., & Heylen, D. (2002). The Virtual Storyteller. *ACM SIGGROUP Bulletin*, 23(2), 20–21.
[Referenced in this thesis from page: 1]
- 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). Association for Computational Linguistics.
[Referenced in this thesis from page: 1, 23, 41]
- Tse, E., Greenberg, S., Shen, C., & Forlines, C. (2007). Multimodal multiplayer tabletop gaming. *Computers in Entertainment (CIE)*, 5(2).
[Referenced in this thesis from page: 16]
- 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 (AVI 2006)*, (pp. 336–343). ACM.
[Referenced in this thesis from page: 16]

Bibliography

- Watson, S., Vannini, N., Davis, M., Woods, S., Hall, M., Hall, L., & Dautenhahn, K. (2007). FearNot! an anti-bullying intervention: Evaluation of an interactive virtual learning environment. In *Proceedings of the Artificial Intelligence and the Simulation of Behavior convention (AISB 2007)*.
[Referenced in this thesis from page: 6]
- Wetherby, A. M. & Wiseman, N. D. (2007). ASD video glossary - glossary of terms. https://www.firstsigns.org/downloads/d_200710_Glossary_of_Terms.pdf.
[Referenced in this thesis from page: 2]
- Whalen, T. (2003). Playing well with others: Applying board game design to tabletop display interfaces. In *ACM Symposium on User Interface Software and Technology*, (pp. 4–5).
[Referenced in this thesis from page: 19]
- Zancanaro, M., Pianesi, F., Stock, O., Venuti, P., Cappelletti, A., Iandolo, G., Prete, M., & Rossi, F. (2007). Children in the museum: an environment for collaborative storytelling. *PEACH-Intelligent Interfaces for Museum Visits*, 165–184.
[Referenced in this thesis from page: 11]

Appendix A

Original interview questions

This appendix contains the original Dutch questions of the informal interviews. The translated English versions can be found in subsection 4.1.1. The answers to the questions are discussed in section 4.5.

Questions after the first interaction session

1. Wat vind je goed aan het systeem dat je zojuist hebt gebruikt?
2. Wat vind je slecht aan het systeem dat je zojuist hebt gebruikt?
3. Wat vind je van de manier waarop je Oma en Roodkapje kan verplaatsen?

Questions after using both versions

1. Wat vind je goed aan het systeem dat je zojuist hebt gebruikt?
2. Wat vind je slecht aan het systeem dat je zojuist hebt gebruikt?
3. (a) Wat vind je van de manier waarop je karakters kan verplaatsen?
(b) In welke versie vond je de manier van karakters verplaatsen beter?
4. Bij welke versie hebben jullie het meeste samengewerkt?
5. Bij welke versie hebben jullie het meeste gepraat?
6. Bij welke versie hebben jullie elkaar het meeste aangekeken?
7. Bij welke versie hebben jullie het meeste gelachen?
8. Welke versie vond je leuker?
9. Welke versie vond je makkelijker?
10. Vond je het leuk om het spel met z'n tweeën te doen?
11. Zou je dit spel ook in je eentje willen doen?
12. Zou je dit spel ook met meer mensen willen doen? Hoeveel?
13. Zou je dit spel nog wel vaker willen spelen? Welke?
14. Heb je wel eens een iPad of andere tablet gebruikt? Hoe vaak?
15. Heb je wel eens een smartphone gebruikt? Hoe vaak?

Appendix B

Quantitative results

In this appendix we present the quantitative data that resulted from the annotations of the first annotator. As mentioned in section 4.3, given the low inter-annotator agreement and the small scale of the evaluation experiment, we could not to draw solid conclusions from the quantitative data. However, for completeness we provide the data and mention some of our cautious presumptions below.

<i>'t-o' = touch-only</i> <i>'t+t' = tangibles+touch</i>	Pair 1		Pair 2		Pair 3		Pair 4	
	t-o	t+t	t+t	t-o	t-o	t+t	t+t	t-o
Net system usage (sec.)	703	542	647	682	829	784	651	876
Wolf thinking (%)	16.1	14.9	22.6	27.3	34.5	33.9	21.4	45.3
Different focus of att.(%)	9.7	14.5	27.7	12.9	9.9	16.1	3.8	6.8
Associative play (%)	54.0	50.0	42.8	38.1	49.3	54.2	43.2	31.4
Social play (%)	21.2	14.8	21.2	31.7	18.5	22.2	17.1	32.0
Highly social play (%)	14.2	18.8	6.3	15.2	19.8	6.8	27.3	20.7
Cooperative play (%)	1.0	2.0	1.7	2.2	2.6	0.8	8.6	9.2

Table B.1: Percentages (of the net system usage time) of the five annotated categories for the first and second run of each pair of children. Additionally, the net system usage time for each run and the percentage of that time where Wolf was thinking are provided.

Second run compared to first run

Pair 1: Second run more ‘higher’ social behaviour

Pair 2: Second run more social behaviour (clearly)

Pair 3: Second run less ‘higher’ social behaviour

Pair 4: Second run more social behaviour? (maybe)

Three times the second run appears to be more social or has more ‘higher’ social behaviour than the first run.

Tangibles+touch compared to touch-only

Pair 1: Tangibles more ‘higher’ social behaviour

Pair 2: Tangibles less social behaviour (clearly)

Pair 3: Tangibles less ‘higher’ social behaviour

Pair 4: Tangibles less social behaviour? (maybe)

Three times the run with tangibles appears to be less social or has less ‘higher’ social behaviour than the first run.

The only time we observed more (highly) social behaviour while using tangibles was with the first pair of children. This outcome could however also be attributed to the fact that this was their second run (in which they might have felt more comfortable).

Influence of having to wait for thinking Wolf

Pair 1: - (No significant difference in thinking time of Wolf)

Pair 2: We see an increase of 4.71 percentage point in Wolf’s thinking time and a clear increase in all kinds of social behaviour.

Pair 3: - (No significant difference in thinking time of Wolf)

Pair 4: We see an increase of 23.97 percentage point in Wolf’s thinking time and an increase in cooperative play and social play, however highly social play slightly decreases.

In both cases where there is a clear difference in the amount of thinking by Wolf, this appears to result in more cooperative play and social play.

Besides the low sample size, one of the biggest problems is to decide how to compare and weigh the results in the three different classes of social behaviour. If we do pursue quantitative analysis in a future experiment, we should specify how these three categories relate to each other. For instance, highly social play is 0.x times as important/desirable as cooperative play and social play is 0.y times as important/desirable as highly social play.