

Interactive Storytelling: A Player Modelling Approach

David Thue¹ and Vadim Bulitko¹ and Marcia Spetch² and Eric Wasylishen¹

¹Department of Computing Science, ²Department of Psychology

University of Alberta

Edmonton, Alberta, Canada, T6G 2E8

{dthue | bulitko | mspetch | ericw}@ualberta.ca

Abstract

In recent years, the fields of Interactive Storytelling and Player Modelling have independently enjoyed increased interest in both academia and the computer games industry. The combination of these technologies, however, remains largely unexplored. In this paper, we present PaSSAGE (Player-Specific Stories via Automatically Generated Events), an interactive storytelling system that uses player modelling to automatically learn a model of the player's preferred style of play, and then uses that model to dynamically select the content of an interactive story. Results from a user study evaluating the entertainment value of adaptive stories created by our system as well as two fixed, pre-authored stories indicate that automatically adapting a story based on learned player preferences can increase the enjoyment of playing a computer role-playing game for certain types of players.

Introduction

In 2006, total revenues for the computer games industry grew to an estimated 30 billion USD worldwide (Plunkett Research, Ltd. 2006). To help ensure that their computer games sell well, game designers have begun to adopt many new technologies in Artificial Intelligence; Player Modelling, the task of learning a player's tendencies through automatic observation, is one such technology. Many recent applications of Player Modelling have sought to improve player enjoyment (and thereby game sales) by automatically adjusting the difficulty of gameplay on the basis of the player's performance (Charles *et al.* 2005). For example, in *SiN Episodes: Emergence*, the likelihood of a fallen enemy dropping a package that heals the player's avatar is determined by previously gathered information concerning the player's shooting accuracy, number of enemy kills, and other factors (Ritual Entertainment 2006). In more general terms, the decision of whether or not the enemy should drop a health pack is made based on a record of the player's past behaviour. What other game decisions might be made using player data?

In this paper, we explore this question in the context of Interactive Storytelling - a story-based experience in which the sequence of events that unfolds is determined while the

player plays. Several varied approaches have been taken toward making decisions in Interactive Storytelling, including chaining together appropriate actor actions, directing scenes toward a dramatic goal, or planning to achieve a learning objective (Crawford 2005; Mateas & Stern 2005; Riedl & Stern 2006). Although these strategies may ultimately improve each player's enjoyment of the game¹, they do so via indirect means. In favour of a more direct approach, we introduce PaSSAGE, an interactive storytelling system which bases its storytelling decisions on an automatically learned model of each player's style of play.

The contributions of this paper are threefold; we: (i) present interactive storytelling as a general decision-making problem, (ii) introduce player modelling as a technique for making storytelling decisions, and (iii) present results from a user study designed to evaluate the proposed technique. The remainder of this paper is organized as follows. We begin by formulating the problem of decision-making in interactive storytelling, and follow with a discussion of related work. We present the details of PaSSAGE's operation, and show and discuss results from a user study designed to test its effect on player enjoyment in a set of short interactive stories. Finally, we suggest directions for future work and draw conclusions concerning the application of player modelling to decision-making in Interactive Storytelling.

Decisions in Interactive Storytelling

Similar to Mateas and Stern's concept of *narrative sequencing*, we present the set of decisions available during storytelling at three levels. At the highest level, decisions must be made concerning the sequence of events that make up the story: How does the story begin? What prompts the protagonist into action? Answering these questions along with many others allows the general structure of a story to take shape. At the next highest level, a more concrete structure is specified by deciding details concerning the time and place of each event, along with the identities of any supporting characters involved. At the third highest level, the behaviours for each character must be determined. We present these levels of decision-making as phases (Selection, Specification, and Refinement, respectively), and discuss each further below.

¹Indeed, Mateas and Stern's *Faade* has met widespread critical acclaim since its release in July 2005.

Event Selection

In most commercial storytelling games, the task of making high-level story-related decisions is complete long before the game reaches store shelves. Traditionally encoded together in tree-like directed graphs, each sequence of events that may be seen by the player is planned out by game authors and fixed prior to release. For the writers of movies and novels, this course is necessary, as only the feedback received before release can be used to adjust their work. In storytelling games, however, this strategy is unnecessarily restrictive; the interactive nature of games provides a mechanism for incorporating post-release feedback in the form of in-game player reactions to story events. Did the player help a man being robbed at the side of the road, join forces with the robbers, or pass by altogether? Each of these potential player reactions contains information on which storytelling decisions can be based, thereby allowing player feedback to be used to refine the story after the game has shipped.

Recent commercial storytelling games use their game environment as a library of smaller directed graphs, where travelling to a new village, country, or planet unveils the starting points for so-called “side-quests” - optional adventures which typically have no significant bearing on a larger story being told (Bethesda Softworks 2006). A common design strategy for such games is to include a wide range of side-quests, each designed to appeal to particular types of players (*e.g.*, dungeon crawls for fighters, treasure hunts for power gamers, etc.), and then rely on players being persistent enough to find the starting points of quests that are suitable for them. Instead of forcing players to sift through a potentially large quantity of uninteresting side-quests to find one that suits their play style, our approach could be used to choose which side-quests should be available at a location before the player arrives, thereby streamlining and improving the player’s experience.

Event Specification

The Event Specification phase aims to answer the question: “*When, where, and for whom should events occur?*” Similar to the choice of story events, these decisions are often made before a game’s release, and technical designers are given all the details necessary to produce the desired scene. Might it be useful to make these decisions post-release as well? If event details such as time, place and the identities of supporting characters (actors) could be decided at run-time, the flexibility and replay value of a storytelling game could be greatly increased; each different decision would provide a new player experience. Our approach implements a method for specifying event details at run-time in a computer role-playing game, as will be presented later on.

Event Refinement

The Event Refinement phase aims to answer the question: “*How should actors behave?*” Unlike the selection and specification of events, the refinement of actor behaviours is regularly performed after the release of current commercial games. Specifically, many games offer varied reactions from actors based on player comments made during dialogue (*e.g.*, rude comments offend, kind comments garner

affection, etc.). Typically, however, the mapping of player comments to actor reactions is fixed long before release; that is, a particular comment (or sequence of comments) will always yield the same reply, for every player. Instead, given the availability of player feedback in the form of in-game actions, the strategy for refining actor behaviours could be decided at run-time. For instance, a player who prefers to avoid conflict should find that actors are less temperamental, while another player who frequently provokes actors might find an argument more often than not. Below we describe an automated way of establishing this mapping at run-time.

Related Work

In 2004, Peinado and Gervás presented a Case-Based Reasoning system designed to mimic the improvisational processes typical of human game masters for pen-and-paper role-playing games (Peinado & Gervás 2004). In such games, one player (designated “game master”) creates an interactive story in which all players play roles, choosing events from his or her imagination to keep all players entertained. Using rules for effective “game-mastering” as well as a set of player types published by Robin Laws, Peinado and Gervás created a system which dynamically selects story events to suit a group of role-players (*e.g.*, if several players indicate that they prefer combat, more combat will occur) (Laws 2001). Unfortunately, players were required to directly identify themselves as preferring one of Laws’ player types; no player modelling was performed.

The event selection mechanism in Mateas and Stern’s *Façade* is based on the concept of *dramatic beats*, a term drawn from the theory of dramatic writing as the smallest unit of dramatic action (Mateas & Stern 2005). During gameplay, beats are selected based on both natural language input from the user on a variety of topics, and an overarching drama manager intended to direct the story along a well-formed Aristotelian tension arc. In 2006, Riedl and Stern used the same technology to create the Automated Story Director, which chooses events according to a partially-ordered plan designed to achieve either a dramatic or educational objective (Riedl & Stern 2006). While Mateas and Stern’s ABL (“A Behaviour Language”) certainly shows promise, they still rely largely on a “one-size-fits-all” approach, as no model of player preferences or play style is maintained.

Recent work by Seif El-Nasr extends Mateas and Stern’s work by incorporating a strategy for player modelling similar to the one presented in this work (Seif El-Nasr 2007). An important distinction between Seif El-Nasr’s system, called *Mirage*, and ours is that hers attempts to model the *player’s character*, while ours attempts to model the *player’s style of playing*. For example, while *Mirage* maintains its model as values along a stereotype of character traits such as heroism, self-interest, and cowardice, PaSSAGE’s stereotype concerns player types (fighter, power gamer, method actor, etc.). Seif El-Nasr’s goal is to allow players to participate in an engaging drama; ours is to maximize each player’s enjoyment of his or her experiences in a virtual world. While maximizing enjoyment may require an engaging drama for some players, others may prefer different kinds of stimulation.

Crawford’s mechanism for selecting “verbs” (Crawford 2005) and Riedl and Stern’s narrative adaptation via “Narrative Directive Behaviours” both offer methods of deferring the final decision of each actor’s behaviour to run-time. As Riedl and Stern point out, one desirable result of doing so is that if an event (or player action) occurs in the story which prevents an actor from believably carrying out its intended behaviour, a different actor may be used in its place. While Crawford’s verb selection is based on actors reacting to their surroundings, our approach most closely follows Riedl and Stern’s technique of proactively searching for actors to carry out desired behaviours (to satisfy a narrative goal).

Player Modelling for Interactive Storytelling

As discussed earlier in the paper, deferring storytelling decisions to run-time can greatly improve the flexibility and replay value of a storytelling game. To achieve this deferral, we introduce PaSSAGE, an interactive storytelling system which uses player modelling to automatically learn the preferred play style of the current player, and uses this knowledge to dynamically adapt the content of an interactive story.

Following Peinado and Gervás, we use the player types from Robin Laws’ rules as the basis for our model; these include Fighters (who prefer combat), Power Gamers (who prefer gaining special items and riches), Tacticians (who prefer thinking creatively), Storytellers (who prefer complex plots) and Method Actors (who prefer to take dramatic actions). During gameplay, PaSSAGE learns a player model expressed as weights for each of these five styles of play; the higher the weight, the stronger the model’s belief that the player prefers that style. Before run-time, potential courses of action are identified by the designer and augmented with weight deltas, allowing the model to be updated based on the player’s actions in-game. For example, the following vectors show how the player model changes when the player asks for a reward in exchange for assistance; since the player is showing an interest in gaining riches, the model’s value for the Power Gamer type increases: (Fighter=1 Method-Actor=81 Storyteller=1 Tactician=1 PowerGamer=41) becomes (F=1 M=81 S=1 T=1 P=141).

Encounter Selection

PaSSAGE tells its stories by drawing from a library of possible events, called *encounters*, each of which has been annotated by an author with information concerning which player types it would be suitable for. For example, being attacked by challenging monsters in a forest might be ideal for players who play as Fighters, and could also appeal to Power Gamers if special items are left behind when the monsters are defeated. Each encounter additionally has one or more *branches* - potential courses of action for the player to take in that situation. When searching for an encounter to run, PaSSAGE examines each encounter’s set of branches, and chooses the encounter whose branch best fits the current values in the player model via an inner-product calculation. To help maintain a strong sense of story, encounters are grouped into sets corresponding to the many phases of Joseph Campbell’s Monomyth (Campbell 1949) - a general structure for

myths that was used prescriptively to create several feature films, including the Star Wars and Matrix trilogies.

Encounter Specification

To make story events independent of time, place, and actor identity, PaSSAGE extends the concept of role passing presented by Riedl and Stern to the game’s environment as a whole (Riedl & Stern 2006); encounters are scripted generically, and details (such as exactly where an encounter should occur) are determined at run-time. This technique helps to eliminate the problem of forcing players to find appealing side-quests, as a new encounter can be chosen based on the model and activated near the player’s current location. For example, consider a section of game-story which requires the player to travel along a road between two cities. Given a set of encounters designed to occur in this setting (alongside a road in the wilderness), the current player model can be used to choose the most appropriate encounter and attempt to activate it between the player’s current position and destination. This activation, however, cannot always occur immediately after selection, as appropriate actors must be obtained to play the roles required by the encounter. In PaSSAGE, we manage this selection via *triggers*: functions which monitor a subset of the game environment, searching for actors which are suitable for the encounter’s roles. Once an encounter has been selected to occur, its trigger function is activated, and the encounter begins once actors have been found to fill all of its roles. In the previous example, a possible trigger function could search for a friendly actor within 20 meters of the player’s current position on the road. As soon as this actor was discovered, an encounter involving a traveller in need would be ready to begin.

Encounter Refinement

Once an encounter and branch have been selected and specified, there is no guarantee that the player will necessarily discover that the chosen branch is a viable course of action. To help alleviate this problem, we introduce a technique called *hinting* to the event refinement phase, similar to Seif El-Nasr’s character improvisations (Seif El-Nasr 2007). Instead of only altering character dialogue as a refinement to each event, we refine the occurrences of the events themselves in an attempt to direct the player toward the chosen branch. For example, for players who play as Power Gamers, dialogue with a demanding character might contain mention of rewards, while Fighters speaking to the same character might find punches being thrown.

Empirical Evaluation

We evaluated PaSSAGE via a user study with respect to the following two hypotheses:

1. **Fun(A) > Fun(F):** Players feel that an adaptive story is more entertaining than a fixed story;
2. **Agency(A) > Agency(F):** Players feel more influential in an adaptive story than in a fixed story.

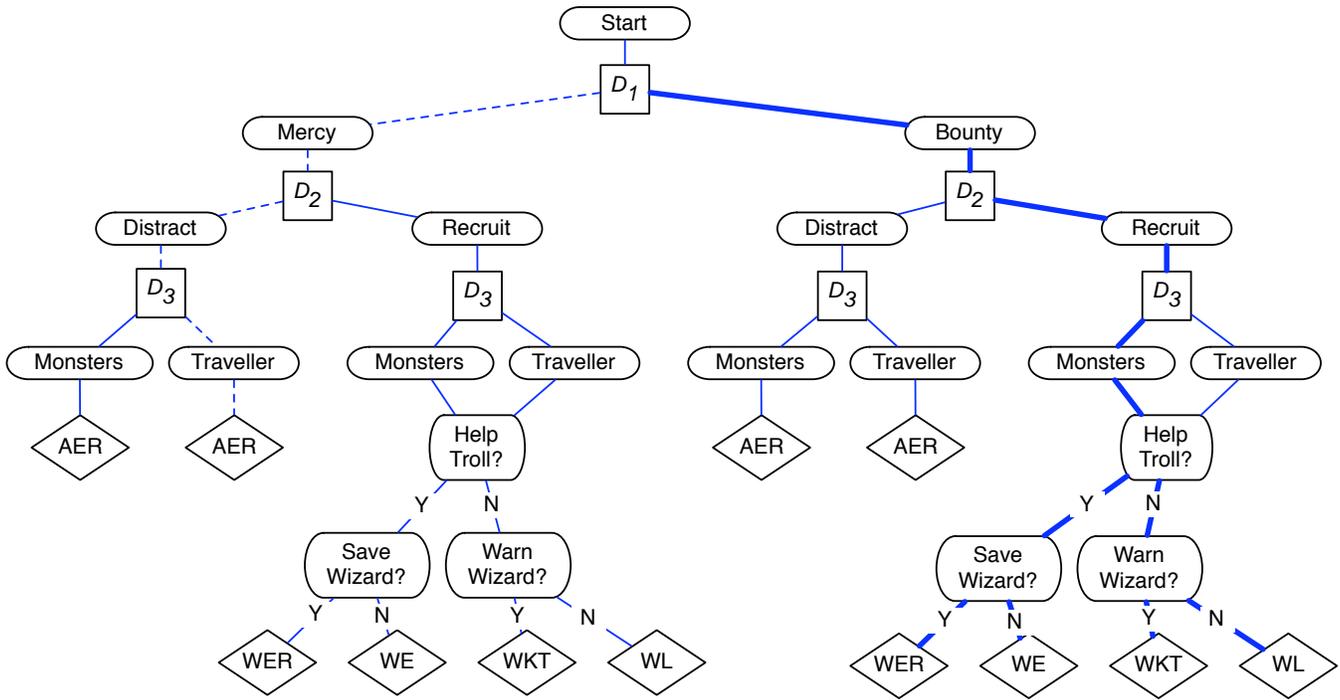


Figure 1: Overview of the game tree. Rounded nodes are in-game events. Diamonds are endings labelled by acronyms: AER = Annara (the player’s character) Eaten & Rescued, WE = Wizard Eaten, WER = Wizard Eaten & Rescued, WKT = Wizard Kills Troll, and WL = Wizard Leaves. Square nodes ($D_{\#}$) represent decisions made by PaSSAGE.

Experimental Setup

Using the Aurora Neverwinter Toolset (BioWare Corp. 2006), we created a library of 8 encounters for a story inspired by the fairy tale “Little Red Riding Hood” (Grimm & Grimm 1812) which recently served as a common ground for discussing techniques in Interactive Storytelling (TIDSE 2006). To avoid a familiarity bias in our study, we changed some of the plot’s elements, all characters, and all dialogue.

We identified three stages of Campbell’s Monomyth in the Little Red Riding Hood story as useful decision points for an interactive storytelling engine: the *Call to Adventure* (Red is sent to Grandma’s house), *Crossing the Threshold* (Red enters the forest and meets the wolf), and the *Road of Trials* (Red faces distractions along the forest path). For each of these stages we authored two potential encounters, each having one or two branches tailored for particular styles of play. To ensure a consistent conclusion to the story, five ending encounters were authored to correspond to the *Ordeal* stage of the Monomyth, and the ending experienced by each player was determined by both previous story events and the player’s immediate actions. Generating all possible sequences of encounters yields the game tree given in Figure 1; it consists of 20 possible lines of gameplay, called *paths*, with five different endings.

Using the game tree as a guide, we created two non-adaptive (“fixed”) stories designed to collectively include every encounter in the tree. Although having one fixed story for every possible combination of encounters would have been ideal, the size of our participant pool limited our tests to only two fixed stories. The first fixed story (shown by dashed

lines in Figure 1) was most closely related to the Little Red Riding Hood fairy tale. The second fixed story (shown in bold lines), led to one of four endings determined directly by player actions. No player modelling was performed during the fixed stories, and all players faced the same decision points (denoted by questions such as “Save the Wizard?”).

In contrast, the adaptive story maintained a player model and had three internal decision nodes (D_1 through D_3 in the figure), each corresponding to one of the three stages of the Monomyth identified above. Each decision node determined which encounter players would face next, based on the model’s current estimate of their preferred style of play. It was possible for players of the adaptive story to traverse any of the 20 paths in the tree, and each player’s path was determined both directly via player decisions and indirectly through the player model.

To help PaSSAGE make an informed decision at node D_1 (*Call to Adventure*), the story began with a “history lesson” in which the player had the opportunity to respond several times to the events of a short sub-story told by an in-game character. Although the player model was updated with each player response, it was not *used* until after the lesson was over. To reduce bias between the adaptive and fixed stories, the history lesson was presented in all three; the player’s responses had no significant effect in the fixed stories.

Adaptive Gameplay Walkthrough

We now follow an actual play of the adaptive story, demonstrating how the model was updated and how it was used in the decision nodes. An annotated video

of this walkthrough can be found at the following URL: <http://ircl.cs.ualberta.ca/games/passage/videos/>

The player finds herself in her room in the basement of her house; the player model begins at initial values (F=1 M=1 S=1 T=1 P=1). Her father arrives and presents a history lesson, during which she has the opportunity to comment several times on the events being related by selecting one of several pre-authored responses from a list (e.g., “*He should have helped the old man - there might have been a big reward!*”). With each comment, the player model is updated. In this case, the player’s choice of comments indicate that her play style is a combination of a strong Method Actor and Power Gamer, and somewhat of a Storyteller. Several model updates occur based the player’s choices, resulting in the following model being created by the end of the lesson: (F=1 M=141 S=41 T=1 P=101). Once the lesson ends, the player is free to wander through the village outside.

As the player exits the house, PaSSAGE’s *Call To Adventure* routine is activated, and an encounter must be selected (D_1). Two encounters are available, both involving the player character’s good friend, Arnell. Given the player’s demonstrated inclination towards being a Power Gamer (P=101), PaSSAGE selects an encounter in which Arnell describes a recently-posted set of bounties available for collection (Bounty in Figure 1) - ingredients for the potion-maker, Jarnas, who lives in the forest. When the player expresses disinterest in collecting the bounties, the model’s Power Gamer value is decreased (F=1 M=141 S=41 T=1 P=61). Although it may have been beneficial at this point to switch to a different *Call to Adventure* encounter, implementing this feature remains as future work; we instead rely on associating a backup motivation with each such encounter. In this case, Arnell relates the mysterious and rare nature of one of the desired ingredients.

When the player enters a nearby forest in search of the potion ingredients, PaSSAGE’s *Crossing the Threshold* routine is called (D_2). Again, two encounters are available, but given the player’s inclination towards being a Method Actor (M=141) and somewhat of a Power Gamer (P=61), an encounter is chosen wherein a troll blocks the player’s path, offering riches beyond imagination for helping to trap an evil wizard in the forest (Figure 2). The player agrees to go along with the dubious plot, and her modelled inclination towards the Method Actor and Storyteller types increase (F=1 M=181 S=81 T=1 P=61). The player continues through the forest while the troll hurries ahead to prepare the trap. Along the way, PaSSAGE’s *Road of Trials* function is called (D_3) to choose an encounter to occur along the forest path. Given the player’s new inclination toward storytelling, the encounter chosen involves a wizard who seems suspiciously similar to a character that was prominent in the history lesson earlier in the day. The wizard asks the player a riddle, which she agrees to solve. The final scene begins with the wizard arriving outside of Jarnas’ house with the troll waiting inside. Instead of luring the wizard into the troll’s trap, the player warns him of the troll’s presence, and a battle between the troll and wizard ensues (Help troll: N, Warn wizard: Y). The wizard slays the troll, and the land is safe once again (WKT in the game tree).



Figure 2: The troll awaiting the player on the bridge.

User Study

To test our hypotheses, we conducted a user study consisting of 90 university students. Once familiar with the interface of our game, each participant played through one of the three stories: two fixed and one adaptive. After finishing the game, they rated their experience along several dimensions including entertainment value, level of interest, replay value, creativity, *etc.* They were also asked to indicate their previous game playing experience, age, and gender.

Table 1 shows statistical significance results for our two hypotheses: (1) that players would find adaptive versions more entertaining than fixed stories, and (2) that players would feel higher agency in adaptive versions; the last two columns give confidence levels for their support. The first two columns represent filters on the participants, designed to highlight segments of the population that might be well-targeted by a commercialization of our approach. A checkmark in the first column indicates that only females (F) are considered. A checkmark in the second column (ETF) limits participants to those who ranked the game as being ‘easy to follow’. A blank in either column indicates no filtering. The columns labelled N_A and N_F list the number of participants for the adaptive and fixed versions respectively.

For example, the first row (✓, ✓) shows that data from females who found the game easy to follow support the hypothesis $\mathbf{Fun(A)} > \mathbf{Fun(F)}$ with a confidence level of 93%. In other words, a T-test with a significance level of 7% ($\alpha = 0.07$) rejects the null-hypothesis $\mathbf{Fun(A)} \leq \mathbf{Fun(F)}$. The last row (two blanks) deals with the data from all participants and fails to strongly support either of the hypotheses.

F	ETF	N_A	N_F	$\mathbf{Fun(A)} > \mathbf{Fun(F)}$	$\mathbf{Agency(A)} > \mathbf{Agency(F)}$
✓	✓	26	33	93%	86%
✓		33	38	85%	84%
	✓	37	39	86%	73%
		45	45	74%	71%

Table 1: Confidence levels in support of our two hypotheses for four data subsets (F = Female, ETF = Easy To Follow).

Another result concerns players of the adaptive version who traversed one of the two fixed paths used in our study (the dashed and bold lines in Figure 1): adaptive-version players who both found the game easy to follow and noted high previous gaming experience on the survey found the game to be more fun (with 80% confidence) than did the fixed-version players of the same paths. All other subgroups we investigated yielded confidence levels below 70%.

Discussion

This first evaluation of our approach has indicated several promising trends. First, females seem to rate PaSSAGE's adaptive stories higher than our fixed stories in terms of both fun and agency. Females who found the game easy to follow support our hypotheses the most, with confidence levels of 93% for fun and 86% for agency. A survey by Lucas and Sherry may offer insight into this result, as they suggest that female players are less likely than male players to have their need for control met in a typical computer game; females are thus more drawn to games in which they experience control, and their enjoyment of games may be more influenced by the degree of agency they experience (Lucas & Sherry 2004). If PaSSAGE's adaptive stories enhanced female players' sense of agency, then finding the game easy to follow might indicate a heightened sense of control, leading to higher scores for both agency and enjoyment. For the adaptive-version plays of the fixed paths, we suspect that our low confidence levels were caused by the study's minimal use of the Specification and Refinement phases of encounter creation.

Future Work

To the best of our knowledge, this is the first controlled evaluation of player modelling for interactive storytelling, and as such, it opens a sizable area for further research. In particular, it will be interesting to explore PaSSAGE's potential with a larger set of encounters, and compare its operation to more fixed paths. While the study presented in this paper focused primarily on validating PaSSAGE in the Encounter Selection phase of story decisions, future studies will focus more intently on the Specification and Refinement phases; we expect that it will be in these phases that adaptive-version plays of fixed paths will show significant improvement over fixed-version plays. Finally, we wish to extend our modelling techniques to include the character traits used in Seif El-Nasr's *Mirage*, as they should be particularly useful for players of the Method Actor type.

Conclusion

This paper made the following three contributions. First, we cast interactive storytelling as a general decision-making problem. Second, we introduced a new system, called PaSSAGE, that automatically maintains a player model and uses it to dynamically select story events during gameplay. Third, we evaluated PaSSAGE via a controlled study of 90 human participants who rated their experience playing fixed or adaptive versions of the classic Little Red Riding Hood story. Compared to the fixed versions, female players who found the game easy to follow felt higher agency in the adaptive version and rated adaptive gameplay as being more fun.

Acknowledgements

We gratefully acknowledge the assistance provided by Jieshan Lu, Marcel Ball, and Cameron Makovichuk in running the user study, and thank all those who either participated in the study or helped beta-test PaSSAGE; their feedback was extremely helpful. We also appreciate support from both the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Informatics Circle of Research Excellence (iCORE).

References

- BioWare Corp. 2006. Aurora Neverwinter Toolset. <http://nwn.bioware.com/>.
- Campbell, J. 1949. *The Hero with a Thousand Faces*. Princeton University Press.
- Charles, D.; Kerr, A.; McNeill, M.; McAlister, M.; Black, M.; Kücklich, J.; Moore, A.; and Stringer, K. 2005. Player-centred game design: Player modelling and adaptive digital games. *Digital Games Research Conf. 2005* 285–298.
- Crawford, C. 2005. *Chris Crawford on Interactive Storytelling*. New Riders.
- Grimm, J., and Grimm, W. 1812. Little red cap. In *Kinder- und Hausmärchen, 1st ed*, volume 1.
- Laws, R. 2001. Robin's laws of good GMing. Steve Jackson Games.
- Lucas, K., and Sherry, J. 2004. Sex differences in video game play: A communication-based explanation. *Communication Research* 31:499–523.
- Mateas, M., and Stern, A. 2005. Procedural authorship: A case-study of the interactive drama *Façade*. In *Digital Arts and Culture (DAC)*.
- Bethesda Softworks. 2006. The Elder Scrolls IV: Oblivion. http://www.elderscrolls.com/games/oblivion_overview.htm.
- Plunkett Research, Ltd. 2006. Entertainment & media industry statistics 2005. <http://www.plunkettresearch.com>.
- Ritual Entertainment. 2006. SiN Episodes: Emergence. <http://www.sinepisodes.com/>.
- Seif El-Nasr, M. 2007. Interaction, narrative, and drama creating an adaptive interactive narrative using performance arts theories. *Interaction Studies* 8(2). To appear.
- TIDSE. 2006. The authoring process in interactive storytelling. Pre-Conference Demo Workshop - 3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment.
- Peinado, F., and Gervás, P. 2004. Transferring game mastering laws to interactive digital storytelling. In *2nd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2004)*, 1–12. Darmstadt, Germany: Springer.
- Riedl, M. O., and Stern, A. 2006. Believable agents and intelligent story adaptation for interactive storytelling. In *3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2006)*, 1–12. Darmstadt, Germany: Springer.