

Socially Consistent Characters in Player-Specific Stories

David Thue¹ and Vadim Bulitko¹ and Marcia Spetch² and Michael Webb¹

Departments of Computer Science¹ and Psychology²

University of Alberta

Edmonton, AB, T6G 2E8, CANADA

dthue | bulitko | mspetch | mrwebb @ ualberta.ca

Abstract

In the context of interactive, virtual experiences, the use of personality models to maintain consistent character behaviour is becoming more widespread in both industry and academia. Most current techniques, however, are limited in one of three ways: either they overly restrict user actions, have a high cost for creating varied content, or rely on a representation that prohibits conveying complex content to the user. Toward addressing these issues, we introduce Socially Consistent Role Passing, a mechanism for ensuring consistent character behaviour that leverages the design of PaSSAGE, an existing system for generating adaptive, interactive stories. While results from previous human user studies have shown that PaSSAGE improves the enjoyment of players with little gaming experience, we present results from a new study showing that PaSSAGE's adaptive stories, augmented with Socially Consistent Role Passing, improve the enjoyment of *all* players versus a set of fixed-structure alternatives.

Introduction

In the commercial video game industry, many games have begun to include models of personality for their artificial characters, toward making them behave more consistently with respect to the player's actions (Bethesda Softworks 2006; BioWare Corp. 2003; 2009). The implicit assumption underlying this inclusion – that consistent character behaviour makes games more enjoyable – seems well motivated, for in the context of any interactive, virtual experience, the consistent behaviour of artificial characters is an important component of the quality of that experience.

The majority of experiences achieve their consistency via hand-crafted scripts, restricting the user's interactions with each character to particular, predetermined scenes. Doing so severely constrains the sets of actions that the user can perform, for future user/character interactions must not be made inconsistent by user actions in the present (e.g., if a character must seek revenge on the user for having been wronged, the user must eventually be forced to act against that character). The common method to avoid forcing or prohibiting user actions is to author alternative sequences of scenes for each character, but the cost of doing so grows exponentially with each new opportunity for the character and user to interact.

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Other, more procedurally-generated experiences ensure consistent character behaviour via sets of author-defined rules (e.g., *Sims 3* (Electronic Arts 2009)). Although such systems succeed in allowing for a variety of user actions, the resulting interactions with characters rarely have the depth of those in the hand-crafted approach above, as their complexity is limited by the ability of current technology to convey the content of their underlying rules. For example, characters in *Sims 3* are able to act maliciously to fulfil a goal of “being mean”, and seem more likely to acquire this goal toward characters who they dislike, but it rests with the user to interpret such actions as being spiteful, vengeful, etc.; there is no explicit way to convey revenge in the game.

In academia, the field of Interactive Storytelling is particularly concerned with the consistent behaviour of artificial characters. In an interactive story, the actions of a user (which primarily include interactions with characters) are used to determine the story's course of events (Mateas and Stern 2003; Barber and Kudenko 2007; Thue et al. 2007a). However, the common approach of driving characters with Artificial Intelligence planners still remains limited in terms of what can feasibly be conveyed to the user, due to the difficulties of generating high quality text from the operators of a planning representation (Thomas and Young 2006).

The challenges described above motivate the creation of a mechanism for maintaining consistent character behaviours which: 1) does not rely on predetermined scenes that are specific to each character, 2) provides more value for each newly created interaction (in terms of story content) than the traditional branching approach, and 3) maintains a high degree of expressive power for authors. To support these requirements, we chose to extend the technique of *role passing* (Thue et al. 2007a), building on our existing system for creating interactive experiences, PaSSAGE (Player-Specific Stories via Automatically Generated Events). In role passing, roles are defined separately for each scene of the experience (called “encounters”) as the behaviours and dialogue that a character should perform if assigned to play that role. Authors constrain the set of possible candidates for each role by specifying a set of properties that acceptable characters must hold (such as being at a specific location in the virtual world or not yet having met the user's character), and characters are assigned to the roles of an encounter at the time that PaSSAGE selects it to occur. PaSSAGE's just-

in-time assignment of characters to roles satisfies point 1 above (no predetermined scenes), and its ability to assign different characters to any given role allows each authored encounter to be instantiated in several different ways; this satisfies point 2 (more value than branching). Given that PaSSAGE uses rule sets only to constrain which actors play which roles (and not the content of the roles themselves), point 3 (maintain expressive power) is satisfied as well. In light of these benefits, our goal was to extend PaSSAGE’s role passing to dynamically ensure that character behaviours are consistent from one encounter to the next.

In this paper, we present Socially Consistent Role Passing, a mechanism for automatically constructing consistent, recurring character roles at run-time, based on each character’s interactions with a user throughout an interactive experience. Following a review of related work, we describe Socially Consistent Role Passing and explain the details of our implementation. We then present the design and results of a human user study that we devised to assess the performance of PaSSAGE with social consistency enabled, and conclude with a discussion of the results and ideas for future work.

Although role passing with social consistency has potential applications in both commercial and educational contexts, the primary aim of PaSSAGE is to create entertaining stories, so we adopt that goal for the remainder of this paper. As such, we will henceforth refer to the interactive experience as a “story”, and its user as “the player”. Furthermore, as the player in PaSSAGE’s stories is always in control of a particular character, treating this character and the player as the same entity will simplify our discussion.

Related Work

In Mateas and Stern’s *Façade* (2003), the story’s two characters each maintain a scalar value that describes how well they like the player, and these values are used to influence the set of content that occurs later in the story, resulting in consistent character behaviour. Although the user’s interaction with each character is not predetermined and the resulting scenes are complex yet conveyed effectively, each element of content required significant author effort to create, and can only be used one time in the experience and by only one particular actor. In Barber and Kudenko’s Generator of Adaptive Dilemma-based Interactive Narratives (2007), each character holds a set of traits (such as attractiveness or morality) which are used to constrain their ability to perform actions in the story and ensure that they behave in a consistent manner. Actions can be performed by more than one actor and the user’s interactions with acting characters are not predetermined, but the presentation of the actions themselves remains relatively shallow (e.g., “You start to fancy Joe” (Barber and Kudenko 2007)). Ochs et al. (2008) present a character personality model similar (though having more dimensions) to that which we describe in this paper. However, instead of using the model to ensure consistent behaviour, they focus on calculating it automatically from the set of emotions that are triggered during interactions with other characters. Rowe et al. (2008) have created a system for generating character dialogue based on characters’ modelled personality and prior interactions with the player, to-

ward improving character consistency with respect to particular archetypical story roles (e.g., “traitor” or “abuser”). Our approach differs from theirs in that while their characters’ personality attributes remain fixed for the duration of the experience, those of our characters are specifically designed to be influenced by the player. Furthermore, we present results of an evaluation of our system, while, to the best of our knowledge, theirs remains to be evaluated. Orkin and Roy (2009) aim to alleviate the burden of authoring new interactions by automatically learning reasonable character behaviours from traces of human behaviour. Although their work with *The Restaurant Game* has shown promising results, the resulting character behaviours are not yet robust enough to satisfy the goals of this work. Mosher and Magerko (2006) have argued that the static aspects of characters’ personalities (e.g., *agreeableness*) are best suited to ensuring consistent behaviour over time. While such traits are certainly useful, we propose that aspects which are more transient (e.g., dynamic character affinity for the player) are also important factors to consider.

Socially Consistent Role Passing

To ensure consistent character behaviours in a story while meeting the three challenges given in the introduction, we take the approach of dynamically constructing recurring roles for characters with whom the player interacts. Inspired by the commercial games cited thus far, we strive to maintain consistency for each character in terms of the *affinity* that it holds toward the player and the actions that it performs. The following sections give details of our approach, including the construction and use of an affinity model in the context of role passing, and a dynamic conversation system designed to highlight consistent character behaviour.

Modelling Affinity

Also referred to as “liking” (Ortony 1991; Ochs, Sabouret, and Corruble 2008), affinity generally describes the degree to which one person likes another: the higher the affinity, the stronger the liking. For simplicity, we chose to model our characters as only having an affinity for the player, and not other (non-player) characters in the story. Each character therefore has a single affinity value which describes how much they like or dislike the player. We represent this value as a continuous variable in the interval $[-1, 1]$, with values near -1 representing strong dislike, zero being neutral, and values near 1 representing strong liking. Updates to the model occur as either positive or negative increments, similarly to Crawford’s B-Numbers (Crawford 2005). Values for these increments (also in $[-1, 1]$) are provided by the authors of the experience as annotations on player actions that relate to other characters. Given an increment to apply, the model is updated using one of the following two equations assuming that $affinity_{max} = -affinity_{min}$.

Affinity update for non-negative increments:

$$affinity = affinity + increment \times (affinity_{max} - affinity)$$

Affinity update for negative increments:

$$affinity = affinity + increment \times (affinity_{max} + affinity)$$

Although some affinity models simply add their increment values to a scalar (e.g., (BioWare Corp. 2009)), doing so makes it awkward for authors to ensure that extreme player actions will have an appropriately large effect on the model. For example, if a particularly heinous act should always result in a negative value of affinity, then subtracting a large negative increment from a scalar is insufficient: the scalar may have grown to be too large to be overcome. Similarly, choosing an increment to force the scalar to some maximal/minimal value is also undesirable, as doing so removes the author’s ability to distinguish between the effects of various extreme actions. With the two equations above, however, the resulting affinity is guaranteed to be negative for all increments in $[-1, -0.5]$, and positive for all increments in $[0.5, 1]$, while increments in $(-0.5, 0.5)$ remain available for less extreme actions.

Using Affinity in Role Passing

Role passing in PaSSAGE works by satisfying a set of author-provided, Boolean constraints on the type of character that should play a given role (e.g., a child within 10 metres of the player). We extended PaSSAGE’s existing set of potential constraints with three constraints pertaining to characters’ current affinity for the player: one which tests affinity against a specified value (e.g., $aff(c) < 0.5$, for some character c), one which tests against the current highest value of affinity held by any character ($aff_{Highest}(c)$: the player’s best friend), and one which tests against the current lowest value held by any character ($aff_{Lowest}(c)$: the player’s worst enemy). These new constraints on affinity allow authors to create roles which are specifically “for” or “against” the player, and ensure that characters will only take on roles which are consistent with their current affinity (e.g., an antagonistic role would only be played by a character having low affinity toward the player).

Recalling Previous Interactions

During role passing, characters with whom the player has interacted at least once are preferred, to help ensure that the player has had an opportunity to influence the affinity that they hold. Given a set of encounters with roles constrained by affinity, particular characters will tend to appear multiple times, behaving each time in a manner that is consistent with their current affinity for the player. We remained concerned, however, that with the existing structure of encounter roles, players may not *remember* having interacted with recurring characters in the past, and therefore not realize that their prior actions were motivating a character’s current behaviour. Toward addressing this concern, we designed the following system for managing character dialogue both inside and outside of encounter roles, with the goal of reminding players of their relevant prior actions.

Dynamic Conversation System

Allowing characters to refer to the player’s prior actions in dialogue necessitates writing brief summaries of the possible outcomes of each interaction. For example, if an interaction with a patron at a bar can end positively or negatively, one

might create two summaries to say “*It was great to see you at the bar!*” and “*I still can’t believe how you treated me at the bar!*” and associate each with the appropriate outcome. We chose to keep these summaries brief (only one sentence long) both to ease their creation, and to avoid overly distracting the player from the *current* interaction that they are meant to support. Whenever such an interaction ends, its outcome is associated with the character involved so that the appropriate summary can later be retrieved.

Summaries can be associated with interactions both inside encounter roles (“role summaries”) and outside of any roles (“chit-chat summaries”). Chit-chat interactions occur when players approach characters who are not actively participating in a current encounter, but are still present to provide a sense of life to the virtual environment (e.g., commoners living in a village). By beginning the experience with a series of chit-chats, players have the opportunity to have an effect on characters’ affinities before any roles are cast for the story’s first encounter.

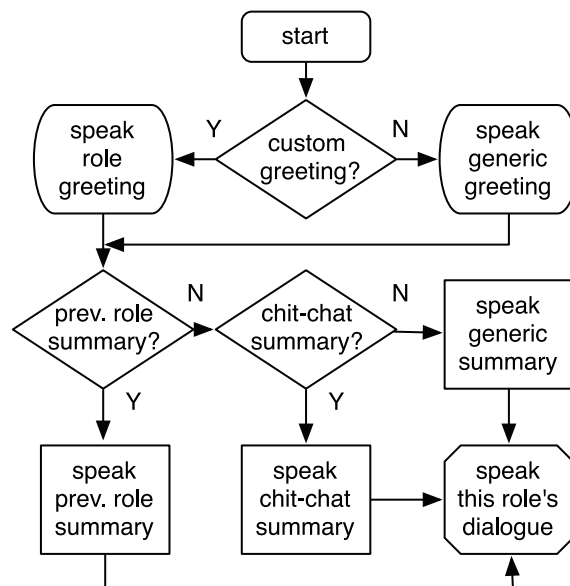


Figure 1: Control flow for the dynamic conversation system.

Once a character has interacted with the player at least once, the dialogue for that character’s next encounter role (should they be cast in one) will be structured as shown in Figure 1. If no specific greeting has been authored for the role, they select one from a library that is consistent with their current affinity for the player; for example, characters with very high affinity will have very positive greetings. Following the greeting, the character speaks the short summary describing their prior interaction with the player (e.g., “*It was great to see you at the bar!*”), preferring to reference summaries from previous roles over summaries from chit-chats. Should the character be required to converse with the player *without* any prior interaction outcomes to summarize, a generic, affinity-consistent summary will be spoken instead (e.g., “*It’s good to see you.*”). Finally, the character continues with the dialogue that was written for their current role (e.g., “*I’ve been hoping to run into you again [...]*”).

Implementation

We have implemented our proposed methods as extensions and modifications to the PaSSAGE source code, using the Aurora Neverwinter Toolset and NWScript (BioWare Corp. 2002); details of PaSSAGE’s use of these tools can be found in our previous publications (Thue et al. 2007a). In brief, PaSSAGE’s stories consist of a sequence of encounters which are selected based on a learned model of the player’s style of play. All players take part in an adventure of three dynamically chosen encounters (for a total of eight possible sequences), followed by one of several endings.

After programming both the model of character affinity and the dynamic conversation system in NWScript, we created a set of chit-chats (with outcome summaries) to help bootstrap the affinity models. These chit-chats occurred during a new scene that we prepended to the story, in which players interact with patrons in a tavern during the night before their adventure begins. All characters were given initial values for their affinity toward the player, which were authored to ensure a fairly even distribution of values (i.e. some characters were initially friendly, while others were initially antagonistic). We substantially modified two of the existing encounters (“Distract” and “Monsters”) to ensure that each had a role that could be constrained by character affinities (see Figure 2). The other encounters required only minor modifications and additions to their role passing constraints (e.g., ensuring that the “Call to Adventure” would be given by a character with high affinity).



Figure 2: Instead of being attacked by spiders themselves, players witness one of their friends (high affinity) being attacked.

Empirical Evaluation

To test PaSSAGE’s effectiveness when augmented with Socially Consistent Role Passing, we conducted a human user study involving 114 undergraduate students (mean age: 19.4 years, 38 were male). Given that PaSSAGE’s primary goal is to provide enjoyable (“fun”) stories, we adopted this metric as the main focus of our study, but also considered players’ reported level of engagement in the story, for we suspected that this metric might be improved in the presence of socially consistent characters.

Our hypotheses for this evaluation concerned three types of story: Adaptive stories generated by PaSSAGE with Socially Consistent Role Passing (A+SCRP), stories with

	A+SCRP	F+SCRP	F
Encounter Selection	Adaptive	Predet.	Predet.
Role Assignment	Consistent	Consistent	Default

Table 1: Differences between story types. A = Adaptive, F = Fixed structure, SCRP = Socially Consistent Role Passing

Fixed structure but with SCRP (F+SCRP), and stories with Fixed structure *without* SCRP (F). Table 1 highlights the differences between each of these types of story. The stories with fixed structure (F+SCRP and F) consisted of specific, predetermined trajectories through PaSSAGE’s story space, with PaSSAGE’s adaptive encounter selection disabled. Stories of type F still contained the same amount of content as those in the other two groups (including the scene with chit-chats at the tavern), but whenever an actor would have been cast into a role by Socially Consistent Role Passing, any affinity constraints were ignored and default Role Passing (i.e., without affinity constraints) was used instead. For conversations, characters always greeted players in group F in a neutral fashion, and no summaries were mentioned. Our hypotheses were as follows:

H1 Players of PaSSAGE’s Adaptive stories with Socially Consistent Role Passing active should have more fun (H1:Fun) and be more engaged (H1:Eng) than players of stories with Fixed structures (predetermined encounters) and Socially Consistent Role Passing active. “(A + SCRP \geq F + SCRP)”

H2 Players of Adaptive stories with SCRP active would have more fun (H2:Fun) and be more engaged (H2:Eng) than players of stories with Fixed structures without SCRP. “(A + SCRP \geq F)”

H3 Players of only stories with Fixed structures would have more fun (H3:Fun) and be more engaged (H3:Eng) when Socially Consistent Role Passing was active versus not. “(F + SCRP \geq F)”

H1 tests PaSSAGE’s adaptive storytelling technology while controlling for SCRP, H2 tests our improvements to PaSSAGE as a complete package, and H3 tests SCRP with story adaptation disabled¹.

Experimental Design

To test our hypotheses, we split our participants into three groups, corresponding to the three types of story given in the previous section. One group played PaSSAGE-Adapted stories with Socially Consistent Role Passing enabled (A+SCRP), one group played stories with Fixed structure but with SCRP enabled (F+SCRP), and one group played stories with Fixed structure *without* SCRP (F).

Like in previous studies of PaSSAGE, all players were told that a group of students had created the story that they were about to play. They were then given a sheet of gameplay instructions to read and allowed a few minutes to practice moving their character around and speaking to charac-

¹Although a full 2×2 design would have allowed for a fourth set of hypotheses (namely, A + SCRP \geq A), we did not collect data for Adaptive story players without SCRP out of concern for splitting our participants into too many small groups and losing statistical power as a result.

ters in the story’s 3D, virtual world (Figure 2). They then played a story corresponding to one of the three groups described above (group assignment was randomized), which all took roughly 25 minutes to complete. Once finished, they filled out a survey indicating how well they enjoyed the experience compared to an average video game of similar length (or their expectation of one), and how engaging they found the experience to be, both on seven point scales with 1 representing negative sentiment (“Less Fun” or “Shallow”) and 7 representing positive sentiment (“More Fun” or “Engaging”). We also asked participants to provide their age, gender, and their average amount of time spent playing video games each week (1: “None at all” to 7: “More than 12 hours per week”).

Balancing Story Trajectories Unlike previous studies of PaSSAGE, which compared the adaptive system against the average results of two particular trajectories through its possible story space (Thue et al. 2007a), we devised a method to guarantee that every possible trajectory through PaSSAGE’s story space would be represented an equal number of times in each of the three study groups. Namely, every trajectory that PaSSAGE chose to fit the player model of one player (group A + SCRCP) was set as a fixed-structure story and experienced by two more players: one with SCRCP active (group F + SCRCP), and one with SCRCP off (group F). Choosing our fixed-structure stories in this way avoided any potential biases that may have been caused by some trajectories being generally more fun or more engaging than others.

Study Results

Table 2 shows the results of running one-tailed t-tests to test each of our hypotheses; each t-test assumes equal variances between the groups being compared, which we verified via Bartlett’s test beforehand. Column 1 shows the metric whose average values are being compared, columns 2 and 3 give the average values (in [1, 7]) for each metric in each group (e.g., the average “Fun” rating for A+SCRCP was 4.11), and column 4 gives confidence values for rejecting the null hypothesis that the value in column 2 is less than or equal to the value in column 3 (corresponding to a one-tailed t-test). Each group consists of data from 38 participants.

Noting that previous versions of PaSSAGE were found to perform particularly well for players with low prior experience with playing video games (Thue et al. 2007b), we also elected to examine this subgroup, defining its members as those who claimed to play at most one hour of video games in an average week (corresponding to 2 or lower on the seven point scale). After balancing over story trajectories for only low experience players, data from 39 participants remained for our analysis (13 per group). Given such a small sample size, we ran Lilliefors’ test to ensure that our samples had normal distributions; we found that the sample corresponding to the metric of “Fun” for group F was not normal (Lilliefors’ p-value = 0.01), and so for the related two comparisons (H2:Fun and H3:Fun) we ran KS-tests for inequality instead. We performed t-tests on all remaining data (with unequal variances for H1:Eng).

H1	A+SCRCP	F+SCRCP	A+S. > F+S.
Fun	4.11	3.61	93.7%
Engaging	4.47	4.16	83.4%

H2	A+SCRCP	F	A+S. > F
Fun	4.11	3.63	93.0%
Engaging	4.47	4.21	79.5%

H3	F+SCRCP	F	F+S. > F
Fun	3.61	3.63	<i>low confidence</i>
Engaging	4.16	4.21	<i>low confidence</i>

Table 2: Mean values and confidence scores for t-tests of six hypotheses. From the top: that (A+SCRCP) > (F+SCRCP), that (A+SCRCP) > (F), and that (F+SCRCP) > (F), for ratings of both “Fun” and “Engaging”. “*low confidence*” is shown in place of scores near 50%, to highlight the fact that they offer very little support for the given hypothesis.

H1	A+SCRCP	F+SCRCP	A+S. > F+S.
Fun	4.46	3.54	94.5%
Engaging	4.92	4.00	96.3%

H2	A+SCRCP	F	A+S. > F
Fun	4.46	3.77	<i>low confidence</i>
Engaging	4.92	4.08	98.3%

H3	F+SCRCP	F	F+S. > F
Fun	3.54	3.77	<i>low confidence</i>
Engaging	4.00	4.08	<i>low confidence</i>

Table 3: Mean values and confidence scores for t-tests of six hypotheses, considering only players with low video game experience (13 players per group). From the top: that (A+SCRCP) > (F+SCRCP), that (A+SCRCP) > (F), and that (F+SCRCP) > (F). H2:Fun and H3:Fun were computed using KS-tests for inequality instead, due to small sample sizes and non-normal data in group F for Fun.

Discussion

Considering the results for all players (Table 2), the confidence values for “Fun” in both H1 (93.7%) and H2 (93.0%) show support for the claim that the adaptive stories generated by PaSSAGE with SCRCP enabled were more enjoyable than their fixed-structure alternatives, both when the latter had SCRCP enabled (H1) *and* when they did not (H2). This confidence score is improved to 94.5% when only players with low experience are considered (Table 3, H1, “Fun”).

For low-experience players, the differences reported for the “Engaging” metric are more pronounced (Table 3, H1: 96.3%, H2: 98.3%), lending some support to our claim that the improved version of PaSSAGE caused (inexperienced) players to feel more engaged than did the fixed-structure stories either with SCRCP enabled (H1), or without SCRCP (H2). As little can be said about the effects of story adaptation on engagement for all players (Table 2, H1 & H2), it may be the case that the novelty of their experience made inexperienced players feel more engaged; this connection between novelty and engagement merits further investigation.

Concerning our hypothesis that considered SCRP on its own (H3), the results are surprisingly inconclusive. The low confidence scores returned by our tests for H3 raise questions concerning the effectiveness of SCRP operating alone on stories with fixed structure, and we lacked the data needed to assess SCRP in the context of adaptive stories alone (A+SCRP vs. A). Is implementing SCRP worthwhile? Perhaps players simply failed to notice that characters were recurring in consistent ways, or there were too few opportunities for the effects of SCRP to be seen. Given the current tendency of commercial story-based games toward including models of affinity for their characters, the value of such efforts should be further explored.

Although the effect of SCRP operating alone remains unknown, its integration with PaSSAGE as a whole represents the most successful version of the software to-date, with confidence levels in our results for all players having increased by nearly 20% (Thue et al. 2007b). Furthermore, given that the same distribution of possible stories was seen by every test group (due to our balancing of story trajectories), we can conclude (with 93.7% confidence) that it was the fact that PaSSAGE adapted its stories to its players that caused them to be more fun.

Future Work

Our extension of PaSSAGE with Socially Consistent Role Passing has yielded several positive results, and opened several interesting questions. The most pressing work that remains is a thorough investigation of the effects of SCRP operating independently of other factors. Smaller scale studies could also be performed, to assess directly whether or not players remember characters from their previous interactions, as well as whether or not such memories have a positive effect on the experience. In terms of the design of SCRP itself, it may be fruitful to add models of mood, relationship, or social dominance alongside the model of affinity. Given such additions, it would be interesting to combine such work with a model of character emotion as Ochs et al. describe (2008). A more formal assessment of the authorial efficiency and expressivity of our approach would also be worthwhile, and future evaluations would likely benefit from the development of a valid survey instrument.

Conclusion

In this paper, we presented Socially Consistent Role Passing, an extension to the PaSSAGE system for interactive storytelling which ensures consistent character behaviours. By dynamically constructing recurring roles based on a model of character affinities for the player, our method remains open to a variety of character actions, offers more usable story content than the traditional branching approach, and retains the ability to present complex story content to its players. We evaluated our method in conjunction with PaSSAGE's existing adaptive storytelling techniques, and found that when compared to a set of fixed-structure alternatives, adaptively generated stories with socially consistent characters are more engaging for players with low gaming experience, and more fun for everyone who played.

Acknowledgements

We appreciate the support of Jim Muirhead with our statistical analysis, and are grateful to the participants in our study. This research was supported by the National Sciences and Engineering Research Council (NSERC) and the Informatics Circle of Research Excellence (iCORE).

References

- Barber, H., and Kudenko, D. 2007. Dynamic generation of dilemma-based interactive narratives. In *AIIDE 2007*, 2–7. Palo Alto, California: AAAI Press.
- BioWare Corp. 2002. The Aurora Neverwinter Toolset. <http://nwn.bioware.com/>.
- Crawford, C. 2005. *Chris Crawford on Interactive Storytelling*. New Riders.
- Mateas, M., and Stern, A. 2003. Façade: An experiment in building a fully-realized interactive drama. *Game Developers Conference, Game Design track*.
- Bethesda Softworks. 2006. The Elder Scrolls IV: Oblivion. http://www.elderscrolls.com/games/oblivion_overview.htm.
- BioWare Corp. 2003. Star Wars: Knights of the Old Republic. <http://www.lucasarts.com/products/swkotor/>.
- BioWare Corp. 2009. Dragon Age: Origins. <http://dragonage.bioware.com/>.
- Electronic Arts. 2009. The Sims 3. <http://www.thesims3.com/>.
- Mosher, R., and Magerko, B. 2006. Personality templates and social hierarchies using stereotypes. In *TIDSE*, volume 4326, 207–218. Darmstadt, Germany: Springer.
- Ochs, M.; Sabouret, N.; and Corruble, V. 2008. Modeling the dynamics of non-player characters' social relations in video games. In *AIIDE 2008*, 90–95. Palo Alto, California: AAAI Press.
- Orkin, J., and Roy, D. 2009. Automatic learning and generation of social behavior from collective human gameplay. In *AAMAS 2009*, 385–392. Richland, SC.
- Ortony, A. 1991. *Memories, Thoughts, and Emotions: Essays in Honor of George Mandler*. Hillsdale, NJ: Lawrence Erlbaum Associates. chapter 22.
- Rowe, J. P.; Ha, E. Y.; and Lester, J. C. 2008. Archetype-driven character dialogue generation for interactive narrative. In *IVA 2008*, vol 5208, 45–58. Tokyo, Japan: Springer.
- Thomas, J. M., and Young, R. M. 2006. Author in the loop: Using mixed-initiative planning to improve interactive narrative. In *Proceedings of the ICAPS-06 Workshop on Preferences And Soft Constraints for Planning*.
- Thue, D.; Bulitko, V.; Spetch, M.; and Wasylishen, E. 2007a. Interactive storytelling: A player modelling approach. In *AIIDE 2007*, 43–48. Palo Alto, California: AAAI Press.
- Thue, D.; Bulitko, V.; Spetch, M.; and Wasylishen, E. 2007b. Learning player preferences to inform delayed authoring. In *Papers from the AAAI Symposium on Intelligent Narrative Technologies*, volume FS-07-05. Arlington, Virginia: AAAI Press. 158–161.