

Learning Player Preferences to Inform Delayed Authoring

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

Of all forms of Intelligent Narrative, interactive narratives are uniquely well-poised to benefit from player modelling techniques. Given the availability of immediate player feedback as interactions with the narrative's world, the traditional task of delayed authoring can be informed with the author's knowledge of which types of players might prefer each event; this allows generated narratives to dynamically adapt and fulfill the player preferences collected by the model. In this paper, we present PaSSAGE (Player-Specific Stories via Automatically Generated Events), our implementation of preference-informed delayed authoring in the setting of interactive entertainment. Recent results from a human user study with 101 participants indicate that for players with minimal previous gaming experience who found the game easy to follow, using our preference-informed techniques can improve their enjoyment of a computer role-playing game.

Introduction

For any commercial venture in Intelligent Narrative to succeed, the resulting product must appeal to a wide audience. Traditional methods for achieving broad appeal typically involve gathering as much audience feedback as possible before release, then striving to reach some compromise between the author's goals and a wide range of audience preferences. Unfortunately, such compromises often result in content which, while appealing to some members of the audience, may greatly displease others (e.g. Half-Life 2's driving sequences, or Façade's social tension) (Valve Corporation 2004; Mateas & Stern 2005). Although this approach is necessary for most forms of narrative, Interactive Narrative allows a mechanism for incorporating audience (player) feedback *after* release: by automatically observing the player's interactions with the narrative's world, her preferences for narrative content can be learned and used by an intelligent narrative system. We refer to this process and its supporting authorial tasks as *preference-informed delayed authoring*; the author describes which types of players would most enjoy each element of narrative content, and his decisions to present or omit certain elements are made at run-time by the narrative system, based on the player's feedback thus far. In this paper, we present PaSSAGE (Player-Specific Stories via Automatically Generated Events), our implementation of preference-informed delayed authoring in the setting of interactive entertainment (Thue *et al.* 2007).

The remainder of this paper is organized into the following three contributions: we (i) present a detailed survey of related methods for informing delayed authoring by monitoring player actions, (ii) extend our previous publication by focusing on our technique for preference-informed delayed authoring in interactive narratives, and (iii) present additional results from our pool of user study participants.

Player-informed Delayed Authoring

As a decision-making process, the authoring of interactive narratives has many opportunities for automation. With intelligent narrative systems, authors encode their ideas of what *might* happen, along with a general knowledge of how things *should* happen. Given this information, a narrative system becomes *informed*; that is, it becomes capable of automatically carrying out authorial decisions at run-time. Although many types of knowledge may be useful in producing a successful narrative (e.g. dramatic theory, narratology), the opportunity for interactive narratives to benefit from player-specific adaptation leads us to focus the following discussion on techniques for gathering and using player feedback as a narrative occurs. We refer to this practice in general as *player-informed delayed authoring*.

In Façade, player feedback is gathered from both natural language utterances and in-game gestures (e.g. handling objects, hugging characters, *etc.*), and recorded as adjustments to a zero-sum affinity value between the game's two non-player characters (NPCs) (Mateas 2002). By presenting the player with events that encourage her to take sides in arguments between the two NPCs, Façade models how well each NPC likes the player. Based on which NPC holds positive affinity for the player, narrative events become either available or unavailable for execution, ultimately helping to determine the conclusion of the story. While this approach to delayed authoring is certainly player-informed, it does not consider which of several available narrative events the player might *prefer* to see happen; ties among available events are broken randomly instead.

In the Interactive Drama Architecture (IDA), a player model is specified in advance as a probability distribution over player actions in the narrative's world (Magerko 2006). Based on this model and an estimation of the player's current knowledge, the likelihood of the player violating the preconditions of pending events is computed and used to ini-

tiate subtle, preemptive corrections when future violations seem likely (e.g. having an NPC burst into laughter in a room nearby to draw the player in). IDA uses its player model to encourage the player to remain within the space of authored story possibilities, but it makes no attempt to identify which of these possibilities its players might prefer.

In *Mirage*, player actions are parsed by a rule-based system to build a profile of the personality being portrayed by the player through her in-game avatar (Seif El-Nasr 2007). *Mirage*'s rules are used to associate a set of author-anticipated player behaviours with adjustments to the personality profile, which is maintained as a vector across five character traits: <reluctant hero, violent, self-interested, coward, truth-seeker>. Like *Façade*, *Mirage* uses its learned player data to affect the selection priority of narrative events. Given an author's specification of how each event's priority should be altered based on the portrayed personality profile, *Mirage* automatically determines the priority of all available events and chooses the event with the highest priority. In addition, *Mirage* tracks the player's cursor movements in an attempt to gauge both her inclination to choose an action and the degree of her hesitation in doing so. Inclinations toward actions are used to prompt behaviour changes in non-player characters (such as blocking the player from leaving a room), and measures of hesitation are used to regulate the magnitude of adjustments made to the personality profile. Although *Mirage* informs several parts of its delayed authoring process with player information, like *Façade*, it remains targeted toward players who appreciate a well-crafted drama. In the following sections, we argue that in interactive entertainment, dramatic tension is only one of the many sources of enjoyment that players seek.

Following David Kiersey's theory of temperaments, Gómez-Gauchía and Peinado attempt to automatically customize the behaviour of their narratives' NPCs using a Case-Based Reasoning approach (Gómez-Gauchía & Peinado 2006). Before the game begins, players fill out a questionnaire; applying Kiersey's theory to the result then indicates their temperament as a proportional combination of four basic types (e.g. <artisan: 20%, guardian: 50%, idealist: 10%, rational: 20%>). Given a knowledge base of dialogue and behaviour variations designed for several player temperaments, Gómez-Gauchía and Peinado's system selects the variation whose associated temperament most closely resembles the current player's. Variations include altering the politeness of the NPC's comments and the speed of the NPC's movements. If the proportions of the selected variation's four basic types fail to match the player's temperament exactly, the system adapts the variation by further adjusting both dialogue politeness and NPC movement speed within restricted ranges. Although the modelling process of Gómez-Gauchía and Peinado's system is independent of the player's actions in-game, one can envision an extension in which a Kiersey-inspired questionnaire could be encoded within the narrative's events, allowing the more seamless technique of learning the player's temperament on-line.

In recent work, Sharma *et al.* extend Nelson and Mateas's work in Declarative Optimization-based Drama Management to incorporate the modelling of player preferences

(Sharma *et al.* 2007; Nelson *et al.* 2006). Through a post-game player survey, ratings were obtained for their overall interest in the presented narrative, their interest in each of the narrative's individual events, and their degree of confidence in their other ratings. By combining the feedback from each play with a trace of its narrative events, Sharma *et al.*'s player model creates a set of cases for a Case-Based Reasoning system, to be used by their drama manager in associating the current player's trace of events with interest ratings for potential future events. Subsequent events are selected via an expectimax search to maximize both a set of predefined author interests and the model's estimation of the player's interests. While Sharma *et al.*'s work does involve informing delayed authoring with player preferences, their dependence on a preexisting case base of player interests limits their flexibility in a commercial setting. Instead of relying on gathering enough pre-release feedback to sufficiently inform the model (or perhaps redistributing post-release feedback via the Internet), we advocate a method of automatically acquiring and using the player's feedback *while* they play, independent of other players' data.

Similar to Sharma *et al.*'s technique of optimizing for player interests, Barber and Kudenko's recent work presents a narrative planner designed to bring about social dilemmas that are expected to be of high interest to the current player (e.g. betraying a friend to achieve great personal gain) (Barber & Kudenko 2007). Set in the context of an interactive soap opera, their system models the player as numerical values across several personality traits: <honesty, faithfulness, selfishness, *etc.*>. The model is updated via author-assigned value adjustments that are attached to each dilemma choice (e.g. +1 honesty, -1 selfishness), and it is used to predict which choice the player will most likely select when presented with a particular dilemma. Given a set of author-defined interest values for each dilemma's choices, the system uses the model's prediction to select the dilemma that maximizes the likelihood of having an interesting outcome for the current player. Barber and Kudenko's work is very near to fitting our definition of preference-informed delayed authoring; while it succeeds in both learning and using player feedback on-line, its use of single, general values for the expected interest of dilemma choices neglects the fact that players with different personalities will likely have different levels of interest in each dilemma choice. Estimating each choice's interest values for various personality models is reflected in part of our approach.

Preference-informed Delayed Authoring

We now present PaSSAGE, our implementation of preference-informed delayed authoring in the context of a computer role-playing game. Following previous work by Peinado and Gervás, PaSSAGE models player preferences as a vector of values for five types of players that Laws identifies as being useful for storytellers in pen and paper role-playing games (e.g. <Fighter (F) = 41, Method Actor (M) = 41, Storyteller (S) = 101, Tactician (T) = 1, Power Gamer (P) = 141>) (Peinado & Gervás 2004; Laws 2001). We begin by detailing the author's tasks in creating a narrative event, and follow by describing how

the player model is learned and used. An example of PaS-SAGE’s operation is presented by Thue *et al.* (2007).

Encounter Creation

In PaSSAGE, all narrative events are *encounters* - events which directly involve the player. As such, each encounter always offers one or more opportunities for the player to act in response; we refer to each such opportunity as a *branch* of the encounter. During the creation of an encounter, in addition to specifying the action that takes place, authors annotate each branch with information concerning what types of players would most enjoy playing along that branch. For example, given an encounter wherein a murderer’s identity is revealed, one possible branch might be to inform the local authorities of his location in exchange for a reward, while a second branch might involve a direct vigilante attack. For the former branch, the author may be writing for Tacticians or Power Gamers, while the latter branch might be preferred by Fighters. Branch annotations are made by encoding values in a vector similar to the one maintained by the player model: *e.g.* $\langle F=0, M=0, S=0, T=4, P=2 \rangle$ annotates a branch that the author expects to be very good for a Tactician and good for a Power Gamer; values of zero for the other types indicate that they are expected to be indifferent toward this branch, and negative values would indicate types for whom this branch should be avoided.

Learning Player Preferences

To learn its model of the player’s preferred styles of play, PaSSAGE must be informed with a set of rules indicating which player types (if any) are indicated by a given player action. These rules can be encounter-dependent (*e.g.* “*if the player demands a reward for locating the murderer, boost the player model’s inclination toward the Power Gamer type (+100)*”), or encounter-independent (*e.g.* “*if the player attacks a friendly NPC, advance the player model along the Fighter type (+40)*”); the magnitude of the adjustment helps distinguish between severe and minor examples. Authors encode all model update rules as the effects of player dialogue choices or world state changes.

Leveraging Player Preferences

PaSSAGE creates its stories as a sequence of encounters drawn from a set of libraries; each library holds encounters designed for a particular phase of Joseph Campbell’s *Monomyth* (Campbell 1949). For each phase, PaSSAGE iterates over the corresponding library, examining the annotations on each encounter’s branches. It calculates the quality of each branch via a clamped inner-product calculation between the branch’s annotation and the values in the player model (negative model values are clamped to zero, as knowledge of the player’s *disinclination* toward certain types is currently unused). The encounter’s selection quality is then taken as the best quality value of its branches. The search for encounters can be initialized with a minimum desired quality or instructed to return the best encounter available.

Empirical Evaluation

In our previous publication, we evaluated PaSSAGE via a user study with respect to the following two hypotheses (Thue *et al.* 2007):

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.

We present additional results in this section using the same experimental setup as before. In brief, we used the Aurora Neverwinter Toolset (BioWare Corp. 2006) to create a library of nine encounters (previously reported as eight by mistake (Thue *et al.* 2007)) based on five stages of Campbell’s *Monomyth*: *Home*, the *Call to Adventure*, *Crossing the Threshold*, the *Trials*, and the *Ordeal*. Each game began with a “history lesson” designed to allow PaSSAGE a chance to model the player before choosing the first encounter, and the adaptive system (with PaSSAGE enabled) was tested alongside two stories with fixed plots chosen to collectively include every created encounter in variations of the fairy tale “Little Red Cap” (Grimm & Grimm 1812).

User Study Extension

To further explore the two hypotheses given above, we extended our user study to 101 participants. Each participant played through one of the three stories (two fixed and one adaptive) and then filled out a post-game survey, rating their enjoyment and sense of influence over the story along scales from 1 to 5. The last two columns of Table 1 give confidence levels in support of our hypotheses. As before, the first two columns represent filters on the participants, designed to highlight segments of the population that might be well-targeted by our approach. A checkmark in the first column indicates that only players who noted having low previous gaming experience are considered (**LE**). A checkmark in the second column (**ETF**) limits participants to those who rated 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 after filtering for both the adaptive and fixed versions, respectively.

For example, the first row (\checkmark, \checkmark) shows that data from players with low prior gaming experience who found the game easy to follow support the hypothesis **Fun(A) > Fun(F)** with a confidence level of 91%. In other words, a T-test with a significance level of 9% ($\alpha = 0.09$) rejects the null-hypothesis **Fun(A) \leq Fun(F)**. As before, the last row (two blanks) deals with the data from all participants and fails to strongly support either of the hypotheses.

LE	ETF	N_A	N_F	Fun(A) > Fun(F)	Agency(A) > Agency(F)
\checkmark	\checkmark	27	30	91%	87%
\checkmark		36	35	80%	88%
	\checkmark	41	45	85%	75%
		50	51	73%	74%

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

Discussion

While the limited length of this paper prohibits a discussion of the larger issues faced by preference-informed delayed authoring, we can instead offer some insight into the results shown above. Our previous publication focused on PaSSAGE's performance for female players; here we presented encouraging data from players with low prior gaming experience. Although our results for players with high prior gaming experience were inconclusive, PaSSAGE's poor performance for such players may have been caused by the following effect: In most commercial story-based video games, when an event occurs that is not directly caused by the player's actions, she assumes (often correctly) that that event had no potential alternatives; this realization detracts from her sense of immersion and enjoyment of the game. However, because PaSSAGE's adaptations were driven indirectly through the player model, experienced players may have assumed (incorrectly) that each selected encounter was the only possible event, rating the game unfavourably as a result. In fact, every encounter after the very first had at least one alternative. Meanwhile, players unfamiliar with interactive narratives in commercial video games may have envisioned a wider range of possible events, making them more appreciative when PaSSAGE chose encounters that were well-matched with their preferred style of play.

Future Work

In our efforts to advance the research of player modelling for interactive narratives, we have temporarily set several concerns aside. The most prominent concern (which is perhaps the most well-treated by others) is the question of what causes one encounter to be in any way causally linked with previous or subsequent encounters. Time constraints led us to solve this problem easily using a branching narrative tree structure, but as others in the field have pointed out, branching narratives are a poor solution at best (Mateas 2002; Magerko 2006). We have taken steps toward moving to a partial-order planning paradigm, and intend to do so as soon as a larger library of encounters has been made. The future work listed in our previous publication also remains.

Conclusion

In this paper, we made the following three contributions. First, we presented a survey of related work in the rapidly growing field of player modelling in interactive narratives. Second, we introduced preference-informed delayed authoring: the task of automatically learning players' preferences by observing their behaviour on-line and using those preferences to dynamically choose the events of an interactive narrative; we presented PaSSAGE as our implementation of this approach in the setting of interactive entertainment. Third, we extended our previous user study to 101 human participants, and found that compared to two fixed narratives, players with low prior gaming experience 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 thank all those who participated in our user study; 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

- Barber, H., and Kudenko, D. 2007. A user model for the generation of dilemma-based interactive narratives. Technical report, AIIDE 2007 Workshop on Optimizing Player Satisfaction, Palo Alto, California: AAAI Press.
- BioWare Corp. 2006. Aurora Neverwinter Toolset. <http://nwn.bioware.com/>.
- Campbell, J. 1949. *The Hero with a Thousand Faces*. Princeton University Press.
- Gómez-Gauchía, H., and Peinado, F. 2006. Automatic customization of non-player characters using players temperament. In *3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE 2006)*, 1–12. Darmstadt, Germany: Springer.
- 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.
- Magerko, B. 2006. *Player Modeling in the Interactive Drama Architecture*. Ph.D. Dissertation, Computer Science and Engineering, University of Michigan.
- Mateas, M., and Stern, A. 2005. Procedural authorship: A case-study of the interactive drama Façade. In *Digital Arts and Culture (DAC)*.
- Mateas, M. 2002. *Interactive Drama, Art, and Artificial Intelligence*. Ph.D. Dissertation, School of Computer Science, Computer Science Department, Carnegie Mellon University.
- Seif El-Nasr, M. 2007. Interaction, narrative, and drama creating an adaptive interactive narrative using performance arts theories. *Interaction Studies* 8(2). To appear.
- Valve Corporation. 2004. Half-life 2. <http://www.half-life2.com/>.
- Nelson, M. J.; Mateas, M.; Roberts, D. L.; and Isbell, C. L. 2006. Declarative optimization-based drama management in interactive fiction. *IEEE Computer Graphics and Applications* 26(3):33–41.
- 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.
- Sharma, M.; Mehta, M.; Ontañón, S.; and Ram, A. 2007. Player modeling evaluation for interactive fiction. Technical report, AIIDE 2007 Workshop on Optimizing Player Satisfaction, Palo Alto, California: AAAI Press.
- Thue, D.; Bulitko, V.; Spetch, M.; and Wasylshen, E. 2007. Interactive storytelling: A player modelling approach. In *AIIDE 2007*. Palo Alto, California: AAAI Press.