Player Agency and the Relevance of Decisions

David Thue[1], Vadim Bulitko[1], Marcia Spetch[2], and Trevon Romanuik[1]

Departments of Computing Science[1] and Psychology[2], University of Alberta
Edmonton, Alberta, Canada, T6G 2E8
{dthue | bulitko | mspetch | trevon} @ ualberta.ca

Abstract. While many forms of storytelling are well-suited to the domain of entertainment, interactive storytelling remains unique in its ability to also afford its audiences a sense of having influence over what will happen next. We propose that the key to encouraging such feelings of agency in interactive stories lies in managing the perceived relevance of the decisions that players make while they play. To this end, we present the design of a system which automatically estimates the relevance of in-game decisions for each particular player, based on a dynamically learned model of their preferences for story content. By actively choosing among several potential consequences of a given player decision, the proposed system highlights the relevance of each decision while accommodating for its players' preferences over potential story content.

1 Introduction

Agency, being the ability to change the course of one's experience, has been found to promote well being in daily life [6], and is a central aspect of video games. Current commercial video games provide a certain type of agency, agency from gameplay, very well; by skillfully manipulating a game's input device, players can explore virtual environments, interact with virtual characters, and overcome challenges as they arise. In the field of Interactive Storytelling, one of the central concerns since its inception has been to provide agency of a different sort: agency from story [10]. Going beyond having players choose their responses to a particular sequence of events, efforts in interactive story generation aim to automatically build a set of many *possible* sequences of events, toward allowing the player's choices to determine which sequence is shown [1, 5, 8, 9].

Is only *providing* players with agency via gameplay or story sufficient to induce strong feelings of control over their experience? Consider the distinction between two phenomena: *theoretical agency*, as one's (objective) ability to change the course of their experience, and *perceived agency*, as one's (subjective) perception of their ability to make such changes. Much research in story generation has focused on finding ways to efficiently provide the former, but always with the implicit assumption that doing so will effectively elicit the latter [5, 8].

In this paper, we challenge this assumption, drawing on research in the field of Social Psychology to suggest that maximizing the perceived agency of players requires more than providing theoretical agency alone; the desirability of the

The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-642-16638-9_26

consequences of player actions must be carefully considered as well. We propose an online algorithm to automatically learn these desires for any given player, and describe a new event-selection routine that dynamically estimates the relevance of each player decision toward increasing the agency of its players.

2 The Control Heuristic

Proposed as a way to model how individuals estimate their degree of control over the occurrence of given outcomes (i.e., perceived agency), the control heuristic was derived by Thompson et al. from an extensive synthesis and unification of prior experimental results [6]. According to the heuristic, one's judgement of control is influenced by two primary factors: their *intention* to achieve the outcome that occurred, and the *connection* that they perceive between their action and its outcome. Intentionality is influenced by three sub-factors: the *foreseeability* of the outcome, one's *ability* to make the outcome occur, and the *desirability* of the outcome for that particular person. If a desirable outcome can be foreseen and one seems capable of achieving it, then intentionality is strongly inferred; the strength of this inference decreases if any of these conditions are not met [6]. Connection is judged in terms of two subtypes: *temporal*, which is stronger the more times a desirable outcome has been observed to occur after an action was taken, and *predictive*, which is stronger when the outcome that occurs was predicted to follow from the action that was taken [6].

Framing Thompson et al.'s work in terms of maximizing the agency perceived by players, it seems that four conditions must necessarily be satisfied for the maximum to be obtained: *foreseeability*, *ability*, *desirability*, and *connection*. These conditions indicate that simply providing players with theoretical agency might *not* be enough to ensure that their agency is effectively perceived: while theoretical agency does grant players the ability to achieve various outcomes and (presumably) also demonstrates temporal connections (*ability* and *connection*), relatively little research to-date has explored how to ensure that the outcomes of player actions are both *foreseeable* and *desirable*. We propose a mechanism designed to address the second of these two conditions, that is, to ensure that the outcomes of player actions are desirable for the particular current player.

3 Related Work

In recent years, an increasing number of commercial video games have granted players the ability to cause multiple story-relevant outcomes [3, 4]. The manner in which they do so, however, seems to be restricted to satisfying players with a particular set of preferences; there is no apparent consideration of which outcomes might be more or less desirable for each specific current player.

In the context of Interactive Storytelling, *GADIN* [1] constructs stories as sequences of difficult player decisions, but its method for selecting subsequent decisions is driven by static, author-defined annotations concerning a general level of interest in each surrounding event. Similarly, *Marlinspike* [9] relies on static annotations concerning the importance of each player action, toward choosing which actions to reincorporate into later story events. Although both systems attempt to provide consequences for player actions, they assume that all players share the same values of interest or importance for each event; we contend that this is not always the case. Mateas has proposed that players perceive agency when a balance exists between two sources of constraints on player actions [2]: the plot of the story (social situations constrain *probable* actions), and the elements of the environment (physical situations constrain *possible* actions). While Mateas' work mirrors the message of the Control Heuristic in terms of *foreseeability* and *ability*, it may yet be improved by considering *desirability* as well. Previous user studies conduced with *PaSSAGE* [7, 8] have shown that automatic story adaptation can improve perceived agency for particular subgroups of players, and we extend that work in the sections that follow.

4 Proposed Approach

Interactive experiences are often conceptualized as a tree of possible states of the environment in which they occur [1, 8, 9], where nodes further from the root occur later during the course of the experience. Nodes represent possible states of the world, and each edge represents the execution of a decision made by either the director, being an Artificial Intelligence system designed to manage the experience (e.g., [5, 8]), or the player. To manage the experience, directors are often able to create, destroy, and change the state of virtual objects in the environment (including virtual characters), as well as initiate story events.

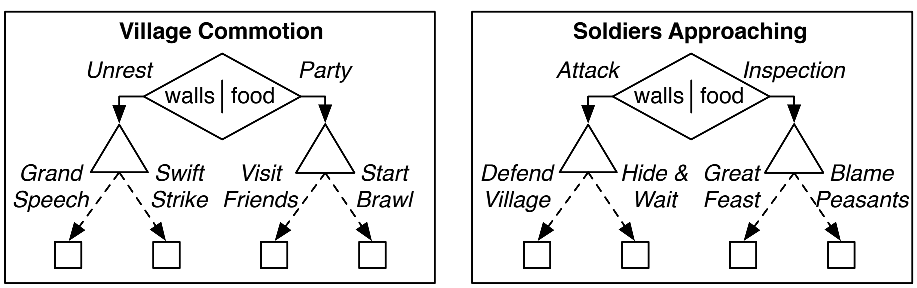

Fig. 1. Two potential events: VC and SA. Triangles are player decisions, dashed arrows are player actions, diamonds are director decisions (based on a prior player decision), solid arrows are director actions, and squares are potential outcomes.

Events in our proposed system are adapted from the *encounters* used in PaSSAGE [8], and can be well thought of as subtrees of the tree that makes up the entire interactive experience. Figure 1 shows examples of two events. Each leaf node of an event's subtree represents an alternate outcome for the event (squares). Edges within an event represent either actions taken by the player (dashed arrows), or actions taken by the director (solid arrows). Players take actions in response to decisions that they are presented with (triangles), and the director takes actions to modify the current event in response to a decision made by the player during some *earlier* event in the story (diamonds). These director decisions divide each event's subtree into two or more components, which we refer to as *sub-events* (*Unrest*, *Party*, *Attack*, and *Inspection*). Labels on arrows describe their actions, and labels in diamonds describe the prior player decision being tested. To illustrate, consider the following hypothetical scenario.

"As the leader of a large village, the player is faced with a dilemma. Rumours abound that bandits in the area have been gathering their numbers to launch an attack, but the village's walls are in a poor state of repair. The walls can be mended for a price, but the state's resources are very low, and the current drought has left the local peasants with very little to eat. Supplemental shipments of food can be ordered from the neighbouring provinces, but the cost of doing so is high. There is enough gold to spend on either the walls or the food, but not both; the player must choose between them."

When choosing an event to occur in the future, the director could have two options, each having two sub-events to select from based on what the player chose to do. For example, one event might involve some commotion in the village that requires investigation, and another might involve reports of soldiers approaching the village walls (VC and SA in Figure 1). If the player chose to repair the village walls instead of feeding the peasants, then the commotion in the village would be the beginnings of a peasant revolt (Director Action (DA): *Unrest*), while the soldiers approaching would be bandits intending to attack (DA: *Attack*). If the player chose to feed the peasants, then the commotion would be the beginnings of a party in the player's honour (DA: *Party*), and the soldiers approaching would be the king of the land arriving to conduct an inspection of the village (DA: *Inspection*). The director's choices of sub-events (*Unrest* versus *Party*, or *Attack* versus *Inspection*) are determined entirely by the player's choice between repairing the village walls or ordering food for the peasants; this design provides all players with a non-trivial amount theoretical agency. The director's decision among *events*, however, is the means by which our proposed system aims to maximize the perceived agency of its players.

4.1 Using Desirability and Relevance to Increase Perceived Agency

Our general strategy is to use the Control Heuristic (Section 2) to improve players' perceived agency. Given its design toward estimating an individual's sense of control over a single outcome, we treat each sub-event as a potential (long-term) outcome of a prior player decision (e.g., walls vs. food in Figure 1). According to the Control Heuristic, perceived agency will be increased when players desire the outcomes (sub-events) that occur as a result of their actions.

Desirability: In *PaSSAGE*, the desirability of potential player *actions* is annotated by authors in terms of inclinations toward five different styles of play:

'Fighter' (F), 'Method Actor' (M), 'Storyteller' (S), 'Tactician' (T), or 'Power Gamer' (P) [8]. *PaSSAGE*'s underlying assumption is that each player can be modelled as a mixture of such inclinations, and that players' enjoyment can be increased by providing them with opportunities to play (i.e., take actions) according to their modelled styles. Taking *PaSSAGE*'s prior success with empirical evaluations as evidence supporting this assumption [7], we have adopted the same player modelling scheme, and so we define desirable sub-events as those which allow players to play in their modelled styles.

Leveraging PaSSAGE's automatically learned player model as well as its calculation of "encounter quality" [7], we propose that the desirability of each sub-event, e, can be calculated based on an inner product between two vectors: the current play-style inclinations in the player model, and the author-provided annotations on each of the sub-event's potential player actions (Equation 1).

$$\mathit{desirability}(e) = \max_{a \in Actions(e)} [\mathit{PlayerModel} \cdot \mathit{Annotations}(a)] \tag{1}$$

To illustrate, for the sub-event *Unrest* in Figure 1, we compute inner products between the player model's vector of inclinations (e.g., [F20,M10,S15,T0,P0]) and each player action's vector of annotations (*Grand Speech*: [F0,M1,S4,T0,P0] and *Swift Strike*: [F4,M0,S0,T2,P0]). Maximizing over the results $(\max(70, 80))$ then provides the desirability of that sub-event ($\mathit{desirability}(\mathit{Unrest}) = 80$).

Relevance: Using the above method to measure the desirability of the outcomes of player decisions, we propose that managing such desirabilities can be well thought of as highlighting the *relevance* of the decisions that players make while playing. Specifically, we define the relevance of a decision as the degree to which it affects the player's ability to experience desirable sub-events during the course of her experience. Given a particular player decision, d, (e.g., 'walls versus food'), we propose that its relevance with respect to a given event, E, can be estimated by the absolute value of the difference between the desirabilities of the sub-event within E that it enables, e^+, and the sub-event within E that it disables, e^- (e.g., in Village Commotion (VC), 'walls' enables *Unrest* and disables *Party*).

$$\mathit{relevance}(d|E) = |\mathit{desirability}(e^+) - \mathit{desirability}(e^-)| \tag{2}$$

Whenever the director must choose a subsequent event, it calculates the relevance of every event that uses d to distinguish between its sub-events (e.g., both events in Figure 1 use $d =$ 'walls versus food'), and chooses the event with highest relevance that also provides a desirable sub-event (i.e., $\mathit{desirability}(e) > 0$). Continuing our example, suppose that the player chose to rebuild the village's walls. The relevance of the two events in Figure 1 would then be computed with respect to this decision, as shown in Equations 3 and 4.

$$\mathit{relevance}(\text{walls}|\text{VC}) = |\mathit{desirability}(\mathit{Unrest}) - \mathit{desirability}(\mathit{Party})| = 50 \tag{3}$$

$$\mathit{relevance}(\text{walls}|\text{SA}) = |\mathit{desirability}(\mathit{Attack}) - \mathit{desirability}(\mathit{Inspection})| = 30 \tag{4}$$

Given that *desirability*(*Unrest*) = 80 (which is positive), the director would choose 'Village Commotion' (VC) as the event which best highlights the relevance of the player's decision, while still providing an outcome that she in particular will desire. According to the Control Heuristic, she should experience increased perceived agency as a result. More formally, we hypothesize that for a given interactive experience having a fixed amount of theoretical agency, proactively choosing events to maximize the relevance of player decisions while providing desirable outcomes will increase the agency that players perceive.

5 Conclusion

In this paper, we posed the question of whether providing players with theoretical agency is sufficient for them to perceive that it exists. Drawing on literature in Social Psychology, we challenged this assumption, and suggested that players must additionally desire the outcomes of their actions to occur. Toward maximizing the agency that players perceive, we presented a method to automatically estimate the desirability of given event outcomes, and proposed the relevance of player decisions as a useful metric for selecting subsequent events. An empirical evaluation via human user study is forthcoming which, if successful, will provide evidence that adapting to players is important for perceived agency to occur.

References

1. Barber, H., Kudenko, D.: Dynamic generation of dilemma-based interactive narratives. In: 3rd Artificial Intelligence and Interactive Digital Entertainment Conference. pp. 2–7. AAAI Press, Palo Alto (2007)
2. Mateas, M.: A preliminary poetics for interactive drama and games. In: Digital Creativity. vol. 12, pp. 140–152. Routledge, London (2001)
3. Bethesda Softworks: Fallout 3. `http://fallout.bethsoft.com/` (2008)
4. BioWare Corp.: Dragon Age: Origins. `http://dragonage.bioware.com/` (2009)
5. Riedl, M.O., Stern, A., Dini, D., Alderman, J.: Dynamic experience management in virtual worlds for entertainment, education, and training. In: Tianfield, H. (ed.) International Transactions on Systems Science and Applications, vol. 3(1), pp. 23–42. SIWN Press, Glasgow (2008)
6. Thompson, S.C., Thomas, C., Armstrong, W.: Illusions of control, underestimations, and accuracy: A control heuristic explanation. In: Psychological Bulletin, vol. 123(2), 143–161. APA Press, Washington (1998)
7. Thue, D.: Player-informed Interactive Storytelling. Master's thesis, University of Alberta, Edmonton (2007)
8. Thue, D., Bulitko, V., Spetch, M., Wasylishen, E.: Interactive storytelling: A player modelling approach. In: 3rd Artificial Intelligence and Interactive Digital Entertainment Conference. pp. 43–48. AAAI Press, Palo Alto (2007)
9. Tomaszewski, Z., Binsted, K.: Demeter: An implementation of the marlinspike interactive drama system. In: AAAI Symposium on Intelligent Narrative Technologies II. pp. 133–136. AAAI Press, Palo Alto (2009)
10. Wardrip-Fruin, N., Mateas, M., Dow, S., Sali, S.: Agency reconsidered. In: Proceedings of the Digital Games Research Association Conference. (2009)