Estimating Player Satisfaction Through the Author's Eyes

David L. Roberts and **Christina R. Strong** and **Charles L. Isbell**
College of Computing
Georgia Institute of Technology
Atlanta, GA 30032
{robertsd,cstrong,isbell}@cc.gatech.edu

Abstract

We explore an approach for inferring player preference functions for interactive entertainment. The goal is to be *adaptive and scalable* by extracting the function through observation, and using a vocabulary that can be understood by game authors rather than depend upon the quirks of individual players. We demonstrate our approach for a number of simulated players in two simulated game environments and present results from an initial feasibility study.

Introduction

There is a growing body of literature on using *drama managers* to guide players through game environments in a narratively consistent way. Drama management has largely focused on the intent of the author rather than the goals of players. In this paper, we consider evaluating satisfaction by *estimating a function of the player's preferences through observed behavior*. We do not characterize a player's behavior exactly, just estimate preferences based on that behavior. There are many types of game players (Bartle 1996); however, recognizing player types and characterizing each type's preferences are orthogonal problems.

Defining the form of the preference function is not straightforward. Players are not typically familiar with formal game analysis and are more likely to describe their experiences in terms of the goals they were able to achieve or the emotional connections they felt with the characters. Different players will also choose different ways to describe their satisfaction. By contrast, game authors are familiar with analysis and game rhetoric. Thus, we turn to the vocabulary of the author as a unifying language.

We make two basic assumptions: (1) players have specific—albeit tacit—preferences that guide their interaction with the game environment, and (2) player satisfaction corresponds to the realization of those preferences. For the purposes of our analysis, we further assume that it is possible to observe the player interacting with a game repeatedly.

Our basic goal is to understand how to model player satisfaction using an author's vocabulary and observations of a player's behavior. If this work is successful, we plan to integrate adaptive preference models into existing drama management techniques. The result will be a drama manager that guides players according to the author's intent but can tailor those experiences to the player's specific preferences.

In the next section, we provide a brief overview of *Declarative Optimization-Based Drama Management*. We then present details of our approach, including player types we consider and an algorithm for estimating player evaluation functions using an author-centric vocabulary. We present the results of an empirical evaluation of our system on a simulated story environment. Finally, we situate this work in the literature and describe the challenges that remain.

Background

Declarative Optimization-based Drama Management (DODM) is a formalism for drama management based on: a set of important plot events with precedence constraints; a set of drama manager actions that influence the game environment and/or the player; a stochastic model of player behavior and a specification of authorial intent in the form of an evaluation function (Nelson & Mateas 2005a; Nelson *et al.* 2006a; 2006b).

An evaluation function encodes the author's story aesthetic. The author simply specifies the criteria used to evaluate a given story, annotates plot points with any necessary information (such as their location or the subplot they advance), and the drama manager tries to guide the story towards one that scores well according to that function. In the process of doing so, it makes complex tradeoffs—difficult for an author to manually specify in advance—among possibly conflicting authorial goals (as specified by components of the evaluation function), taking into account the player's actions and incorporating them into the developing story.

Generally speaking, there is a common vocabulary for defining authorial goals using a small set of story features that can be used to describe most narrative experiences from the author's perspective. To make weighting various authorial goals straightforward, all features range from 0.0 to 1.0, so an author can specify an overall evaluation function as a weighted combination of the features. Seven features have been studied in earlier work (Nelson *et al.* 2006a).

Location flow is a measure of spatial locality of action: The more pairs of plot points that occur in the same location, the higher the score. This feature is based on a judgment that wandering constantly around the world is undesirable.

Thought flow measures continuity of the player's (assumed) thoughts, as specified by an optional *thought* annotation on plot points. This feature prefers very short snippets of coherent "sub-subplots"; for example, *get_safe_combo* and *discover_safe* are both annotated with the thought *safe*, so the thought-flow feature would prefer plots in which the player finds the safe and then looks for the combination (or vice versa), rather than finding the safe, getting distracted by something else, and then finding the combination later.

Motivation measures whether plot points happened *apropos* of nothing, or happened after other motivating plot points. For example, first finding the observatory and noticing that the telescope is missing a lens would make opening the puzzle box and finding a lens well-motivated, while opening the puzzle box without having found the observatory would make the discovery of the lens un-motivated.

Plot mixing measures how much the initial part of the story includes plot points from multiple subplots. One might want the player to explore the world in the beginning, rather than finding one of the plot sequences and going straight to one of the endings.

Plot homing is a counterpart to plot mixing. It measures to what extent the latter part of the story includes plot points from the same subplot. While we may not want the player to move directly to one subplot and finish the game right away, we probably do want her to eventually follow a coherent story, rather than continually oscillating between subplots and then stumbling upon one of the endings.

Choices is a measure of how much freedom the player has to affect what the next plot point will be. The goal is to allow the player as many choices of action at any given time as possible, rather than achieving a highly-rated story simply by forcing the player into one. Without this feature, a drama manager might linearize the story, making the best story as judged by the other features the *only* possible story, defeating the purpose of an interactive experience. This feature can be seen as a way of trading off just how much guidance the drama manager should give the player.

Manipulativity is a measure of how obvious the drama manager's changes in the world are. The author specifies a manipulativity score for each DM action, encoding a judgment of how likely that action is to be noticed by the player as something driven by the manager. A hint to go through a door (*e.g.*, by having the player hear someone talking in the next room) might be judged less manipulative than forcing the player to enter a door (*e.g.*, by locking all other doors).

Although these features sometime refer to the state of mind of the player, they have not been used to describe directly a story from a player's perspective. Part of our goal is to determine if these features can be used to estimate a player's preference function.

Our Approach

In this section, we describe our approach to characterizing player satisfaction. Recall that we make two assumptions. First, we assume that players have preferences that guide their interaction with the environment (even though they may not be able to articulate those preferences). Second, we assume that the more preferences the player satisfies, the more enjoyment they derived from the experience.

Eliciting preferences from humans can be very difficult, so we instead try to observe player preferences. We have a language for describing story quality in terms of story features, so we attempt to construct an estimate by learning weights of a linear evaluation function over those features. Other models are possible, but a linear function is simple. We leave more complicated models for future investigation.

Algorithm 1 details this approach. All features, weights, and functions associated with the author will be given an "a" superscript; all features, weights, and functions, associated with the player will be given a "p" superscript; and all estimated weights and functions will be marked with a "~". For example, an estimated player evaluation function defined over author features would be represented as: $\tilde{e}^p(t) = \sum_{i=1}^{F} \tilde{w}_i^p \cdot f_i^a(t)$ (as in Line 5 of Algorithm 1).

Algorithm 1 Comparing actual to estimated player evaluations and to authors evaluations.

1: Fix a player evaluation function $e^p(t) = \sum_{i=1}^{F} w_i^p \cdot f_i^p(t)$.
2: Convert the evaluation function $e^p(t)$ to a player model $P(t'|a,t)$ by considering locally greedy behavior.
3: Sample a set of complete game traces T^p without the use of the drama manager.
4: Convert each trace in T^p to its feature vector representation.
5: Estimate the weights of a player evaluation function defined over author features $\tilde{e}^p(t) = \sum_{i=1}^{F} \tilde{w}_i^p \cdot f_i^a(t)$ using the stories in T^p.
6: Plot a histogram of story evaluation quality to frequency for every $t \in T^p$ using $e^a(t)$, $e^p(t)$, and $\tilde{e}^p(t)$.
7: Learn a DM policy for $P(t'|a,t)$ and sample a set T of complete game traces using it.
8: Plot a second histogram of story evaluation quality to frequency for every $t \in T$ using $e^a(t)$, $e^p(t)$, and $\tilde{e}^p(t)$.

To run our experiments, we first choose a player type and use it to construct a player model (details of this process are presented in the next section). Once the player model is fixed, it is used to sample a set of stories without a drama manager. Because the player is allowed to act on her own without any influence from a drama manager, we assume that this set of sampled stories is a representation of how the player wishes to behave in the game. These stories are then converted to a feature representation and the frequency of each instantiation of a feature vector is used as the measure of its desirability. In other words, we assume that the more often a certain combination of features occurs, the more desirable it is to the player. Thus, we use the feature vector instantiations and frequencies as input for a regression problem. Then, a DM policy is learned for the specific player model and we sample another set of stories using the DM to determine how well the player would be satisfied under the DM that optimally represents the author's intent.[1]

[1] We solve for this stochastic policy using *Targeted Trajectory Distribution MDPs* (TTD-MDPs). A complete discussion of TTD-MDPs is well beyond the scope of this paper. The interested reader is referred to (Roberts *et al.* 2006; Bhat *et al.* 2007).

Player Models and Evaluation Functions

In this work, we consider three basic player types: (1) the fully cooperative player who has the same evaluation function as the author; (2) the partially cooperative player who shares the same set of features with the author, but has her own set of weights; and (3) the independent player who has her own set of features and weights.

The independent player has many subtypes: an explorer who likes to visit as much of the story world as possible; a non-explorer who does not; a habitual player who tends to prefer the same subset of plot points strongly; a social player who prefers plot points centered around other characters; a non-social player who does not; and a player who likes to accumulate objects and prefers plot points based on objects.

Each independent player type defines a set of weights over player-specific story features. Some of these player specific features are similar to those of the author. For example, the author's *Choices* feature can be adapted to the player that enjoys exploring the game world. We have selected a set of features for this evaluation that we feel are useful in describing the behavior of some intuitive player types. It is not intended to be an exhaustive set of features. Yannakakis and Hallman have worked on measuring player satisfaction by observing and extracting features of their behavior external to the game (like heart rate) (Yannakakis & Hallam 2007). In contrast to that work, we use features that are internal to the game. Here, we define the player specific features that do not correlate well with the author's features.

Location is a measure of the number of unique locations a player has visited. This allows explorers to discover as many new places as possible.

Social Interaction measures the amount of interaction a player has with non-player characters.

Habits indicate that a player prefers a specific set of plot points. This is annotated in the plot points themselves.

Object Discovery measures the number of objects a player has discovered. Objects include ones the player can pick up and those permanently fixed in a specific location.

Note that these features are not mutually exclusive even for our basic players. For example, the social player may still apply weight to Object Discovery because having certain object may provoke discussion with a character. As mentioned above, this set of features is not intended to be exhaustive; however, it is intended as a starting point by which we can judge the feasibility of this approach.

In order to turn the player evaluation function into a player model, $P(t'|a,t)$, we assume the player acts in a greedy way. At every decision point, the player evaluates the partial story that consists of the plot points encountered thus far and one of the possible next plot points to obtain a score. Scores are then converted to a distribution and the player chooses the subsequent plot point according to that distribution.

Estimation

As noted above, we are interested in estimating the weights of a linear evaluation function using a set of stories. Unfortunately there are a super-exponential number of possible stories. Even with thousands of samples, we are unlikely to see each story more than once or twice.

Recall that the author defines a set of features for describing stories. Therefore, we can consider a story t to be represented by a vector $\vec{v}(t) = [f_1^p(t), f_2^p(t), \ldots, f_F^p(t)]$. To overcome the sparsity problem, we opt to consider stories as a vector of features rather than as trajectories of plot points. Thus, multiple stories are represented by the same feature vector, giving us more apparent samples.

In practice, this is not a complete solution. We found that there were still a relatively large number of feature vectors that appeared only a few times. Thus, we further abstracted a story into a "binned" feature vector. Specifically, if we want b bins and the evaluation of a particular feature is $f_i^p(t)$, then the binned value is $\frac{\lfloor f_i^p(t) \cdot b \rfloor}{b}$. For the experiments presented below, we use 10 bins. The system of equations is:

$$\begin{bmatrix} r(\vec{v_1}) \\ r(\vec{v_2}) \\ \cdots \\ r(\vec{v_n}) \end{bmatrix} = \begin{bmatrix} f_1^a(t_1) & f_2^a(t_1) & \cdots & f_F^a(t_1) \\ f_1^a(t_2) & f_2^a(t_2) & \cdots & f_F^a(t_2) \\ \cdots & \cdots & \cdots & \cdots \\ f_1^a(t_n) & f_2^a(t_n) & \cdots & f_F^a(t_n) \end{bmatrix} \cdot \begin{bmatrix} w_1^p \\ w_2^p \\ \cdots \\ w_n^p \end{bmatrix} \quad (1)$$

where $r(\vec{v_i})$ is the number of occurrences of feature vector $\vec{v_i}$, $f_i^a(t_j)$ is the feature evaluation of a story t_j that produces feature vector $\vec{v_j}$, and w_i^p is the weight of feature i. Rewriting Equation 1 as $\vec{R} = \vec{F} \cdot \vec{W}$ the best weights are:

$$\widetilde{W} = (\vec{F}^T \vec{F})^{-1} \vec{F}^T \vec{R}$$

Characterizing Success

In existing work using this drama management formalism, one presents results as histograms over story quality (Weyhrauch 1997; Lamstein & Mateas 2004; Nelson & Mateas 2005b; 2005a; Nelson *et al.* 2006a; 2006b). A set of stories are evaluated according to the author's evaluation function and the results are plotted with the evaluation result as the independent variable and the frequency of evaluation as the dependent variable. Comparing the histogram of non-drama managed stories to the histogram of drama managed stories should result in a shift up and to the right.

Here, we would like similar results; however, we are not necessarily looking for the same positive shift. Instead, we seek to show that the change in the shape of the histogram for the estimated evaluation function mirrors that of the change in shape of the player's actual evaluation, regardless of the direction of that change.

Results

We conducted a number of experiments on two different story worlds. First, we considered the simulated story *Alphabet City* originally studied by Roberts *et al.* (2006). Second, we examined an abstraction of a subset of the text-based interactive fiction *Anchorhead*. For each of the experiments, we used 5,000 stories for evaluation. Additionally, the results we present are an average over 10 trials.

First we consider experiments conducted on Alphabet City where the player shares the same set of features as the author. In Figure 1, we plot the quality distribution for a player that is completely cooperative with the author (*e.g.* has the same exact evaluation function). We have added an offset to the player's curves to distinguish them from the author's curves as their shapes are identical. Notice how in

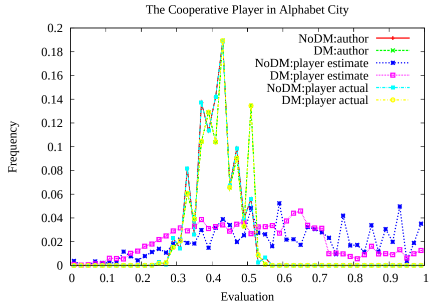

Figure 1: Comparison of the author's, player's, and estimated quality distribution for a player with the same evaluation function as the author in Alphabet City. Note that an offset has been added between the player and author curves to distinguish them.

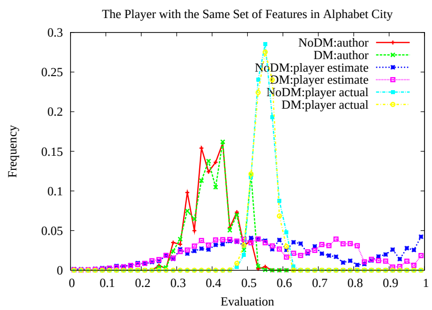

Figure 2: Comparison of the author's, player's, and estimated quality distribution for a player with the same set of features but a different set of weights than the author on Alphabet City.

this case, there is a clear difference in the estimated curves: the curve with drama management is noticeably lower than the curve without drama management in the high end of the evaluation range and vice verse in the low end. Unfortunately, this points out that even when the player shares the evaluation function of the author, the estimate is not particularly accurate; however, the magnitude of the difference between the curves is significant. As we shall see below, this is a recurring aspect of this model.

In Figure 2, we present results for a player with the same set of features as the author, but a different set of weights. In this case, the distribution according to the player's actual evaluation bears little resemblance to the author's; additionally, notice that, although less pronounced than in Figure 1, a similar "positive" shift is observable in the distribution of author evaluations with and without the drama manager. On the other hand, the change between the player's distribution of evaluations with and without drama management is almost imperceptible. In this experiment, we found that the difference between the distribution over estimated player quality with and without the drama manager is actually significantly closer than before. This is encouraging because the player's actual evaluation changes very little as well.

Next, we consider the results of experiments conducted on Anchorhead. Figure 3 shows the histograms for experiments using the non-exploring player. Note that the use of the drama manager in this case actually produces a "negative shift" in the distributions according to the author's evaluation function. This behavior is not unexpected. It can be explained by the "falling off" phenomenon discussed by Roberts *et al.* (2006). It is also interesting to note that the relative shape of the distributions according to the author and player are very similar despite the fact that their evaluation functions are defined over a different set of features.

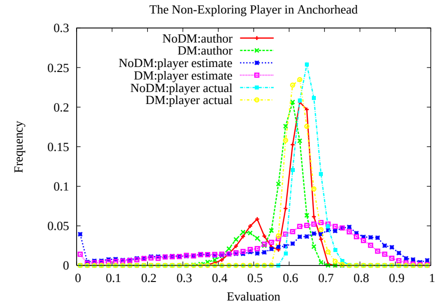

Figure 3: Comparison of the author's, player's, and estimated quality distribution for the non-exploring player on the Anchorhead "god" subplot.

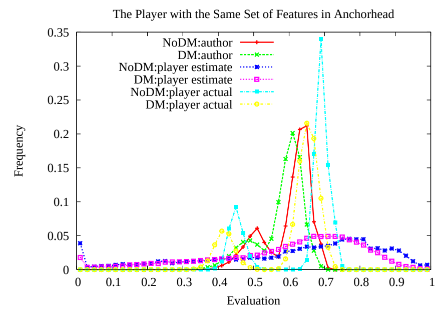

Figure 4: Comparison of the author's, player's, and estimated quality distribution for a player with the same features but different weights than the author on the Anchorhead "god" subplot.

Figure 3 alone is interesting; however, considered together with Figure 4 where the results of the experiments conducted using the player with the same set of features in Anchorhead are presented, we begin to see some interesting behavior. First, notice in Figure 4 that the shape and relative position of the author's curves are similar. Additionally, in contrast to the previous figure, notice that the difference in shape and position of the player's actual evaluation curves is much more significant than in previous experiments. Bearing that in mind, compare the shape and relative positions of the estimated curves with and without drama management in Figures 3 & 4. In Figure 3, where the difference between the player curves is less noticeable, the difference in the estimated curves is also slightly less noticeable. On the other hand, in Figure 4, where the difference between the player curves is more noticeable, the difference between the estimated curves is also slightly more noticeable.

To explore this relationship further, we compare the estimated curves for the non-exploring player and the player

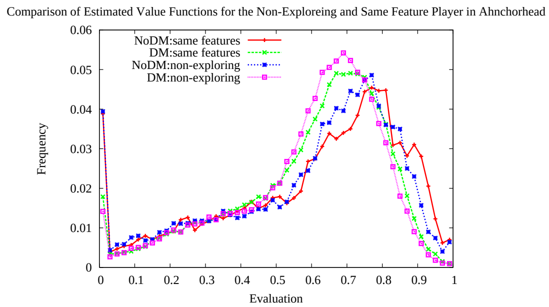

Figure 5: Comparison of the estimated quality distribution for the non-exploring player and the player with the same set of features but a different set of weights than the author in Anchorhead.

with the same features in Anchorhead in Figure 5. This is a more detailed comparison of the estimated curves in Figures 3 & 4. As noted before, the difference between the player's curves is more pronounced for the player with the same set of features than it is for the non-exploring player. In comparing the associated estimated curves in this figure, we see that this more pronounced difference is also observable in the estimated curves.

Related Work

Space does not permit a thorough review of the growing body of literature on drama management. We direct the reader to the cited work on DODM or TTD-MDPs, and to Mateas (1999) and Roberts & Isbell (2007) for surveys.

With respect to incorporating player preferences in a drama management system, there has been some recent efforts, such as that of Sharma *et al.* (2007). In that work, a case-based reasoning approach[2] is used to model "player preference." The authors set their model apart from the models of "player behavior" that are the basis of work on DODM and TTD-MDPs. During game play, they identify relevant preference models in their "case base" by considering the sequence of plot points that have occurred. Each of these models has a preference score that is used to characterize quality from the player's perspective. This preference score is obtained through actual player evaluation performed after an episode of game play. Additionally, the quality of the match between the current player and the model from the case base is used to skew the drama manager's decision to include both authorial intent and player evaluation.

This approach to drama management appears to be the only one that explicitly targets player preference. Although promising, it has several drawbacks. Accurate elicitation of player preference through questions can be tricky at best. Further, player preference may be non-stationary and non-transferable (*e.g.*, the player may change her preferences across episodes and one player's preferences may not accurately model another's). Lastly, evaluating the approach is difficult. If a player reports a good experience, it is difficult to tell if the cause is the drama manager, the author having specified a good narrative, or the player just being overly nice. In our work, we seek to avoid this complication by providing a computational approach to evaluating quality.

As for evaluation, there has been little work done with respect to drama management. One approach independent of a drama manager is that of Sweetser & Wyeth (2005). They propose a framework for modeling player satisfaction in games that is based on "game flow." Their model has eight core features from game rhetoric literature that are intended to characterize player experiences: concentration, challenge, player skills, control, clear goals, feedback, immersion, and social interaction. While this model is compatible with our approach to evaluation, it is designed specifically for game analysis. In other words, it is designed to be a subjective criteria for experts to use as an evaluation framework for games. In that sense, it is a predictive model of player satisfaction, not a descriptive model as we seek to build. Still, the features used in that work could be used to estimate a player evaluation function instead of the author-defined features we used here. It remains to be seen how to provide a computational representation of some of their features, such as player skills. In addition to this work, Malone has proposed a set of heuristics for designing "fun" instructional games (1980).

In her thesis, Federoff (2002) provides a broad survey of other approaches to designing game experiences to maximize player satisfaction. She classifies the literature into three categories: (1) interface controls and display; (2) game mechanics and interaction; and (3) game play problems and challenges. Each of these categories is focused on system design, and not authorial quality. In this work, as in that of Sweetser & Wyeth, we seek to describe the satisfaction the player derives from the narrative quality of the game.

Discussion and Future Work

One potential problem with this work is the need to have a large number of samples to perform regression accurately. This can be especially difficult when these samples must be obtained from an actual human player. Fortunately, there are some strategies for addressing these concerns. First, we would like to conduct experiments to determine the sensitivity of this approach to the number of sampled stories used as input for regression. Additionally, we believe we can further reduce this burden by considering the local decisions of the player; that is, we could look at the individual plot points, rather than stories, and try to estimate a function of their feature values based on the frequency of occurrences. As it is a simpler model, it requires fewer sampled stories.

Using frequency as the target for regression reveals that the use of the DM steers the player away from their normal habits. Unfortunately, this approach does not allow us to characterize the quality of that shift. In the case where the player has the same evaluation function as the author, one can see a large shift in quality; however, the direction of that shift appears to be opposite of what we would expect because any influence exerted by the DM will appear to be negative in terms of the estimated curve, as it will change the frequency of feature occurrences in a set of stories. In

[2] Case-based reasoning is a "lazy" method for function approximation. It stores previously seen examples and uses those to construct estimates for new ones. The interested reader is directed to Kolodner & Leake's tutorial for a more thorough discussion (1996).

future work, we plan to examine techniques that will allow us to overcome this limitation. Specifically, rather than look at the frequency of complete stories, we plan to leverage information from local decisions. For example, if the player frequently takes transitions to plot points that preserve location flow, we can try to learn a model of the player that weights that feature more heavily.

As noted before, player preferences can be non-stationary. For example, once the player has fully explored the game environment, she may hone in on certain parts of the game experience that she liked. Thus, any attempt to take user behavior into account—especially across repeated play—must use some form of online adaption. We believe that considering local decisions will be important in realizing this goal.

In addition, once we have honed our approach and are better able to mirror the qualitative shift in the player's actual evaluation function using our estimated evaluation function, we plan to incorporate this into the drama manager's decision making. Considering a combination of the author's evaluation and the player's evaluation in fixing the target distribution for TTD-MDPs will allow a policy to be learned that can make the trade-off between authorial intent and the player's autonomy to pursue her own goals.

Finally, we intend to run a series of user studies to validate the assumptions we have made about player preferences and player satisfaction.

Acknowledgments

This research was performed while on appointment as a U.S. Department of Homeland Security (DHS) Fellow under the DHS Scholarship and Fellowship Program, a program administered by the Oak Ridge Institute for Science and Education (ORISE) for DHS through an interagency agreement with the U.S Department of Energy (DOE). ORISE is managed by Oak Ridge Associated Universities under DOE contract number DE-AC05-06OR23100. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DHS, DOE, or ORISE. We also acknowledge the support of DARPA under contract No. HR0011-06-1-0021.

References

Bartle, R. 1996. Hearts, clubs, diamonds, spades: Players who suit MUDS. *Journal of MUD Research [online]* 1(1). http://www.mud.co.uk/richard/hcds.htm.

Bhat, S.; Roberts, D. L.; Nelson, M. J.; Isbell, C. L.; and Mateas, M. 2007. A globally optimal online algorithm for ttd-mdps. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS07)*.

Federoff, M. 2002. *Heuristics and usability guidelines for the creation and evaluation of fun in video games*. Ph.D. Dissertation, Indiana University, Bloomington. Unpublished thesis, http://www.melissafederoff.com/thesis.html.

Kolodner, J. L., and Leake, D. B. 1996. A tutorial introduction to case-based reasoning. In Leake, D. B., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*.

Lamstein, A., and Mateas, M. 2004. A search-based drama manager. In *Proceedings of the AAAI-04 Workshop on Challenges in Game AI*.

Malone, T. W. 1980. What makes things fun to learn? heuristics for designing instructional computer games. In *SIGSMALL '80: Proceedings of the 3rd ACM SIGSMALL symposium and the first SIGPC symposium on Small systems*, 162–169. New York, NY, USA: ACM Press.

Mateas, M. 1999. An Oz-centric review of interactive drama and believable agents. In Woodridge, M., and Veloso, M., eds., *AI Today: Recent Trends and Developments. Lecture Notes in AI 1600*. Berlin, NY: Springer.

Nelson, M. J., and Mateas, M. 2005a. Search-based drama management in the interactive fiction Anchorhead. In *Proceedings of the First Annual Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-05)*.

Nelson, M. J., and Mateas, M. 2005b. Search-based drama management in the interactive fiction Anchorhead. In *Proceedings of the First Annual Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-05)*.

Nelson, M. J.; Mateas, M.; Roberts, D. L.; and Isbell, C. L. 2006a. Declarative drama management in the interactive fiction anchorhead. *IEEE Computer Graphics and Applications (Special Issue on Interactive Narrative)* 26(3):30–39.

Nelson, M. J.; Roberts, D. L.; Isbell, C. L.; and Mateas, M. 2006b. Reinforcement learning for declarative optimization-based drama management. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-06)*.

Roberts, D. L., and Isbell, C. L. 2007. Desiderata for managers of interactive experiences: A survey of recent advances in drama management. In *Proceedings of the First Workshop on Agent-Based Systems for Human Learning and Entertainment (ABSHLE 07)*.

Roberts, D. L.; Nelson, M. J.; Isbell, C. L.; Mateas, M.; and Littman, M. L. 2006. Targetting specific distributions of trajecotries in MDPs. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)*.

Sharma, M.; Ontañón, S.; Strong, C.; Mehta, M.; and Ram, A. 2007. Towards player preference modeling for drama management in interactive stories. In *Proceedings of the Twentieth International FLAIRS Conference (FLAIRS07)*.

Sweetser, P., and Wyeth, P. 2005. Gameflow: A model for evaluating player enjoyment in games. *ACM Computers in Entertainment* 3(3):1–24.

Weyhrauch, P. 1997. *Guiding Interactive Drama*. Ph.D. Dissertation, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA. Technical Report CMU-CS-97-109.

Yannakakis, G. N., and Hallam, J. 2007. Game and player feature selection for entertainment capture. In *Proceedings of the 2007 IEEE Symposium on Computational Intelligence and Games (CIG 2007)*.