

Using Feature Value Distributions to Estimate Player Satisfaction Through an 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 preferences for interactive entertainment. The goal is to be *adaptive and scalable* by extracting preferences through observation, and using a vocabulary 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 games and present results from a 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. Research has largely focused on the intent of the author rather than the goals of players. Here, we consider evaluating player satisfaction by *estimating a function of the player's preferences*. The goal is not to characterize a player exactly—for example by identifying one of the many types of players she might be—but to estimate her preferences *based only on her behavior*.

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, many 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. These assumptions may seem strong; however, consider Bartle's description of the four player types (Bartle 1996). Each of the four player types has certain characteristics of their game play style. The “killer” for example enjoys the disruption of other player's game experience. Therefore, in realizing this preference for disruption, the killer will likely obtain the most satisfaction.

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 this paper, we extend previous work in which the estimate of the player's evaluation function was constructed by looking at the distribution of stories as represented by feature vectors (Roberts, Strong, & Isbell 2007). In that work, each occurrence of a story was represented as a binned feature vector based on the set of features the author supplies for evaluation. In contrast, here we consider the distribution of feature values on a per-feature basis. As we shall see later, this approach yields notably more accurate results in most cases.

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 and compare the results to those presented in earlier work. Finally, we situate this work in the literature and describe open challenges.

Background

Declarative Optimization-based Drama Management (DODM) formalizes for drama management based on: a set of important plot events with precedence constraints; a set of drama manager actions that influence the game experience; a stochastic model of player behavior; and an evaluation function specifying authorial intent (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.

Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Generally speaking, there is a common vocabulary for defining authorial goals using a small set of story features. For simplicity, 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 constantly wandering 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 early, rather than finding a plot sequence that moves quickly to an ending.

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 reward finishing the game right away, we probably do want the player to eventually follow a coherent story, rather than oscillate between subplots before 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. 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 measures how obvious the drama manager is. The author specifies a score for each DM action, encoding 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 through a door (*e.g.*, by locking 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 to what extent features can be used to estimate a player’s preference function.

Our Approach

Recall our two assumptions. First, we assume that players have preferences that guide their interactions. Second, we

assume that the more preferences the player satisfies, the more enjoyment they derive from the experience.

While humans have preferences, they may not be able to articulate them, so eliciting them directly can be difficult. Instead we infer preferences by observing the player’s behavior. We have a language for describing story quality in terms of story features, so we construct an estimate by learning a function over those features. Many models are possible, but we begin with a simple linear function. We leave more complicated models for future investigation.

For our experiments, we first sample a set of stories using a player model that acts uniformly random and no drama manager. Next, we choose a specific player type and use it to construct a player model. Once the player model is fixed, it is used to sample another 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. Both of these sets of stories are then converted to a representation of feature distribution based on the frequency of each instantiation of a feature’s value. The important point here is that *as the difference between the uniform baseline and unmanaged player feature value distributions grows, the importance of that particular feature in characterizing a player’s behavior grows as well*. We measure the importance of each feature using one of two ways of comparing distributions, and use our measure to construct an estimate of the player’s evaluation function. Then, a DM policy is learned for the specific player model and we sample another set of stories using the DM. Lastly, the set of unmanaged stories are compared to the set of stories under the DM that optimally represents the author’s intent using the player’s actual evaluation function as well as the learned evaluation function.¹

Before detailing our approach, we introduce some notation. 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 “ $\hat{\cdot}$ ”. For example, an estimated player evaluation function defined over author features would be represented as: $\tilde{e}^p(t) = \sum_{i=1}^F \hat{w}_i^p \cdot f_i^a(t)$.

Let N be the number of histogram bins and let \mathcal{T} be a set of stories. Then, we define the distribution $F_i^a(\mathcal{T})$ of the author’s feature f_i^a as a histogram where the value of the b -th bin is given by:

$$F_i^a(b, \mathcal{T}) = \frac{\sum_{\forall t \in \mathcal{T}} \begin{cases} 1 & \text{if } \frac{b}{N} \leq f_i^a(t) < \frac{b+1}{N} \\ 0 & \text{otherwise} \end{cases}}{|\mathcal{T}|} \quad (1)$$

where $|\mathcal{T}|$ is the number of stories in \mathcal{T} .

We define two comparison measures: *Kullback-Liebler (KL) divergence* and *mean shift*. The *KL*-divergence be-

¹We solve for this stochastic policy using *Targeted Trajectory Distribution MDPs* (TTD-MDPs). The interested reader is referred to (Roberts *et al.* 2006; Bhat *et al.* 2007).

tween distributions $p(t)$ and $q(t)$ is:

$$\begin{aligned} D_{KL}(p||q) &= \sum_t p(t) \log \frac{p(t)}{q(t)} \quad (2) \\ &= \sum_t p(t) \log p(t) - \sum_t p(t) \log q(t) \quad (3) \end{aligned}$$

KL -divergence is not a true distance, as it is asymmetric; however, it is a well-understood measure with several important properties. In particular, it is consistent, always non-negative and zero only when p and q are equal. On the other hand, we define $D_\mu(p||q) = \mathbf{E}[p(t)] - \mathbf{E}[q(t)]$ to be the mean shift between two probability distributions.

Player Models and Evaluation Functions

We consider three player types: (1) the fully cooperative player with the same evaluation function as the author; (2) the partially cooperative player who shares the set of features with the author, but has different weights; and (3) the independent player with 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 measures the number of unique locations a player has visited. Explorers seek to maximize this value.

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 objects may provoke discussion with a character. Some of these features represent player types discussed by Bartle (Bartle 1996).

In order to turn the player evaluation function into a player model, 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

Even with thousands of samples, we are unlikely to see each story more than once or twice. This can make estimation difficult. 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 $\vec{v}(t) = [f_1^a(t), f_2^a(t), \dots, f_F^a(t)]$. To overcome the issue of sparsity of examples, we opt to consider stories as vectors 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 $\lfloor \frac{f_i^p(t) \cdot b}{b} \rfloor$. For the experiments presented below, we use 10 bins.

In earlier work (Roberts, Strong, & Isbell 2007), we considered the frequency of feature vectors (or stories) as the measure of story goodness. Under this assumption, the estimated player evaluation function can be obtained by performing linear regression. In that case, the system of equations is:

$$\begin{bmatrix} r(\vec{v}_1) \\ r(\vec{v}_2) \\ \vdots \\ r(\vec{v}_n) \end{bmatrix} = \begin{bmatrix} f_1^a(t_1) & f_2^a(t_1) & \dots & f_F^a(t_1) \\ f_1^a(t_2) & f_2^a(t_2) & \dots & f_F^a(t_2) \\ \vdots & \vdots & \ddots & \vdots \\ f_1^a(t_n) & f_2^a(t_n) & \dots & f_F^a(t_n) \end{bmatrix} \cdot \begin{bmatrix} w_1^p \\ w_2^p \\ \vdots \\ w_n^p \end{bmatrix} \quad (4)$$

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 4 as $\vec{R} = \vec{F} \cdot \vec{W}$ the best weights are:

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

The elements of \vec{W} are the weights for the player evaluation function in terms of the author’s features.

This approach did not perform particularly well. In contrast, by considering distributions over individual feature values (e.g. the value of $f_i^a(t) \forall t \in \mathcal{T}$ rather than the distribution over feature vector instantiations), we obtain an estimate of *feature importance* rather than *story importance*. Focusing on feature importance allows us to better capture the way in which a player acts differently than uniformly randomly, especially under sparse data.

Thus, we construct an estimate of the player’s evaluation function in the following manner. Let \mathcal{T}^u be a set of stories sampled with a uniformly random player model and let \mathcal{T}^p be a set of stories sampled with the actual player model.² Then, using the notation introduced above, we define the estimate of the player’s evaluation function as:

$$\tilde{e}^p(t) = \sum_{i=1}^N \left(\frac{|D(F_i^a(\mathcal{T}^u)||F_i^a(\mathcal{T}^p))|}{\mathcal{Z}} \cdot Q(f_i^a(t)) \right) \quad (5)$$

²In practice, this set of stories could be observed from actual human players.

where

$$Q(f_i^a(t)) = \begin{cases} f_i^a(t) & \text{if } D_\mu(F_i^a(\mathcal{T}^u) || F_i^a(\mathcal{T}^p)) \geq 0 \\ 1 - f_i^a(t) & \text{otherwise} \end{cases} \quad (6)$$

D in Equation 5 is either D_{KL} or D_μ depending on the particular experiment, $\mathcal{Z} = \sum_{i=1}^N D(F_i^a(\mathcal{T}^u) || F_i^a(\mathcal{T}^p))$ is the normalizing constant, and $Q(f_i^a(t))$ is a “direction” function that captures the sign of the difference between the distributions’ means. We will also report on a number of experiments conducted “without direction.” In those cases, $Q(f_i^a(t)) = f_i^a(t)$.

Thus, we have four types of experiments: mean shift weights both with and without direction and KL -divergence weights both with and without direction. Note that regardless of the method used to determine weights, the mean shift is used to determine direction. Thus, we use mean shift to determine if a difference indicates the player has a desire to realize high or low values for that particular feature.

Characterizing Success

In existing work using this drama management formalism, results are presented 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 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. In successfully capturing the shape change between the baseline curve and the curve using a DM, we can make qualitative claims as to the satisfaction of the player. In particular, we can claim that the player gained satisfaction in part because of changes made by the DM.

Results

We conducted a number of experiments on three 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*. Third, we looked at a re-implementation of *Tea for Three*, originally studied by Weyhrauch (1997). For each of the experiments, we used 5,000 stories for evaluation. Lack of space prevents us from presenting detailed results for all of the experiments we ran; however, we will show graphs for representative results and also provide a summary of all of the experiments.

Representative Results

First, we will highlight the vast improvement our new feature-based approach has over the older story-based ap-




Figure 1: Comparison of the cooperative player’s actual evaluation and estimated evaluation function for Alphabet City. In this plot, the estimated evaluation function was constructed using the previous story distribution approach.

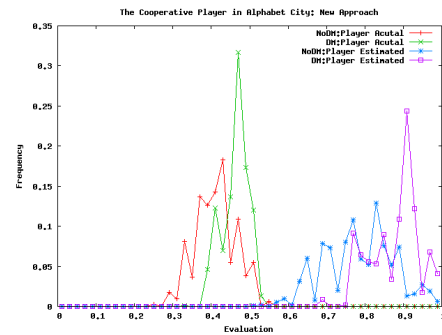


Figure 2: Comparison of the cooperative player’s actual evaluation and estimated evaluation function for Alphabet City. In this plot, the estimated evaluation function was constructed using the new feature distribution approach, with KL -divergence weights and the Q direction function.

proach. Consider Figure 1, where one of the most successful trials using the story-based approach is presented for the cooperative player (*e.g.* has the same evaluation function) in Alphabet City. Notice how in 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 versa 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. This is a recurring result of the story-based approach.

Now, consider Figure 2. These results represent the same player and story but use the feature-based approach. In this experiment, the weights of the estimated evaluation function were obtained from KL -divergence and the Q direction function was used. Compare the relationship of the NoDM curve to the DM curve for both the actual and estimated functions. Note how all four curves are (essentially) uni-modal, and that for the estimated and actual cases, the NoDM curves are higher than the DM curves toward the lower end of the range and lower than the DM curves to-

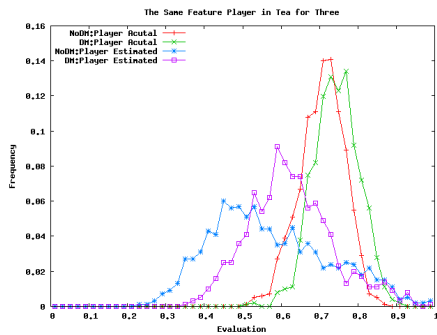


Figure 3: Comparison of the same featured player’s actual evaluation and estimated evaluation function in Tea for Three using the new approach (KL -divergence weights and using the Q direction function).

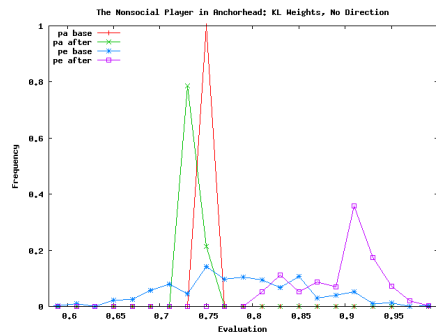


Figure 5: Comparison of the nonsocial player’s actual and estimated evaluation functions in Alphabet City using KL -divergence without the Q direction function.

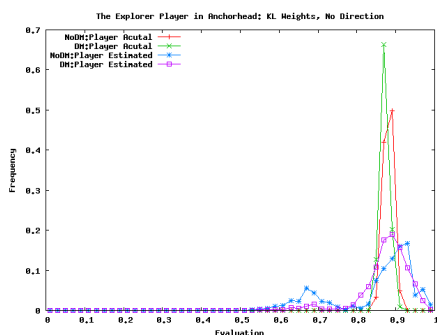


Figure 4: Comparison of the same featured player’s actual evaluation and estimated evaluation function in Anchorhead using the new approach (KL -divergence weights but without using the Q direction function).

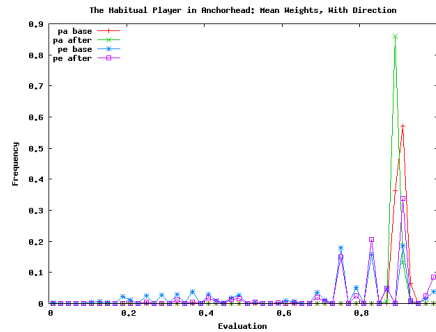


Figure 6: Comparison of the habitual player’s actual evaluation and estimated evaluation function in Anchorhead using mean shift weights but using the Q direction function.

ward the higher end of the range. This is exactly the behavior we are looking for. Matching the actual curve exactly is not what is important; rather, we want to mirror the qualitative shift between the NoDM and DM curves apparent in the actual evaluation curves using the estimated curves. In our summary table, this result would be considered *excellent*.

Now consider Figures 3 & 4. They show results for the player with the same features (but different weights) as the author in Tea for Three and in Anchorhead. In Figure 3 we used KL weights and direction whereas in Figure 4 we used KL weights but no direction. In Figure 3, the results are almost exactly what we would hope for. Notice how qualitatively the shape change between the NoDM and DM curves is similar for both the actual and estimated curves; however, notice that at the highest part of the range, the estimated curves are basically similar whereas the actual curves are not. We will refer to such a result as *good*. Next, consider Figure 4. The results are less convincing. Notice how the actual set of curves is uni-modal but the estimated curves are bi-modal. In the actual case, the DM curve is slightly higher and left of the NoDM curve. On the other hand, while the estimated curve has a similar shape in the higher of the two modes, the lower mode has clear separation between

the NoDM and DM curves. We refer to this sort of result as *mediocre*.

Next, we consider an example of *backwards* results. In Figure 5 we present the results of the nonsocial player in Anchorhead using KL -divergence weights and not using the Q direction function to construct the estimate. The effect of the DM on the actual curve is to move the mass of the distribution down and left (generally negative) whereas the effect of the DM in the estimated curve is to move the mass up and right.

Lastly, we present in Figure 6 an example of *bad* results. This experiment was run on Anchorhead using the habitual player and the estimate was constructed with mean shift weights and the Q direction function. First, note how there is no basis for a structural comparison between the actual and estimated curves—the actual curves are nicely uni-modal whereas the estimated curves are highly kurtotic. Note also that there is a noticeable negative qualitative shift between the two actual curves. Further, it is very difficult to characterize the qualitative shift between the estimated curves.

Summary of all Results

We examined three different stories and used four methods for estimating the weights of the player evaluation function. As noted, we classify our results into five categories: *excel-*

lent, good, mediocre, backwards, and bad. Table 1 contains results for all of our experiments.

First, we point out that experiments run using mean-shift weights did not yield good results. Using mean shift weights with Alphabet City and Anchorhead results in estimated curves that do not reflect the qualitative shift of the actual curves in any way. In the case of Tea for Three, the results were slightly better. Most trials resulted in backwards curves that at least capture some of the structure of the underlying problem. We believe this shows that the mean is an insufficient statistic for capturing the qualitative shift we care about.

Second, we note that the experiments run using KL -divergence weights were much better. Perhaps even more encouraging, a majority of the bad or backwards results obtained using KL -divergence weights came from experiments run on Alphabet City. In fact, almost none of the experiments run on Alphabet City yielded positive results. We believe this is because of the relatively low number of author features defined over this story. The authors of Alphabet City defined their evaluation function using only four features whereas the authors of Anchorhead and Tea for Three defined seven. We believe that the increased representational power obtained with the extra features enables a more accurate characterization of the player’s behavior.

Third, note the experiments run with KL weights and either Anchorhead or Tea for Three. 24 out of these 32 experiments (75%) yielded results that would enable an accurate claim as to the change in player satisfaction. Additionally, looking only at the experiments run on those stories using KL weights and the Q direction function, we see 100% success in obtaining results that are at least somewhat reflective (*i.e.* mediocre or better) of the actual qualitative shift.

Fourth, we found that the quality of the results was very high for the cooperative and same-featured players. This is somewhat unsurprising, but indicates that the approach can handle the cases that should be easier.

Finally, we also noticed that certain player types were captured better by the author’s features than others. For example, the object discovery player and the habitual player tended to produce worse results than the explorer and non-explorer. Unfortunately, our experiments were not designed to test the robustness of our approach to variations in parameters such as the number of author features. As such, making a definitive claim about which player types are better represented using author features is difficult. In the future, we plan to conduct experiments to look at the interaction between the player features and author features more closely.

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 approach³ 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.

In addition, there have been a number of efforts to construct player models. For example, Magerko utilizes a “predictive” player model in his Interactive Drama Architecture to reason about threats to the narrative progression that player actions might cause (2006). In contrast, Seif El-Nasr uses a “descriptive” model based on parameters that represent different player types (2007). Based on a series of if-then rules, the parameter values are updated according to player behavior in the game. This is somewhat similar to the player model used by Thue *et al.* (2007; 2006) and stereotype-based models (Rich 1979; Tsiriga & Virvou 2002). Others have taken approaches such as estimated players’ goals and plans (Kautz & Allen 1986) or estimated players’ mood and emotions (Picard 2003). One thing that all of these systems have in common is that they build models online. Our approach is offline. We build our model from player traces obtained beforehand and use those models during an episode to affect the drama manager’s decision making.

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

³Case-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).

	mean, no direction			mean, direction		
	<i>Alphabet</i>	<i>Anchorhead</i>	<i>Teafor3</i>	<i>Alphabet</i>	<i>Anchorhead</i>	<i>Teafor3</i>
social	bad	bad	backwards	bad	bad	backwards
nonsocial	bad	bad	backwards	bad	bad	backwards
explorer	bad	bad	backwards	bad	bad	backwards
nonexplorer	bad	bad	backwards	bad	bad	backwards
objectdiscovery	bad	bad	backwards	bad	bad	bad
habitual	bad	bad	backwards	bad	bad	bad
same	bad	bad	excellent	bad	bad	mediocre
author	bad	bad	excellent	bad	bad	backwards
	<i>KL, no direction</i>			<i>KL, direction</i>		
	<i>Alphabet</i>	<i>Anchorhead</i>	<i>Teafor3</i>	<i>Alphabet</i>	<i>Anchorhead</i>	<i>Teafor3</i>
social	backwards	backwards	backwards	backwards	good	mediocre
nonsocial	backwards	excellent	backwards	backwards	excellent	mediocre
explorer	backwards	mediocre	backwards	backwards	excellent	excellent
nonexplorer	backwards	mediocre	backwards	backwards	mediocre	excellent
objectdiscovery	backwards	bad	backwards	backwards	good	backwards
habitual	backwards	good	backwards	backwards	mediocre	backwards
same	good	excellent	excellent	good	excellent	good
author	good	excellent	excellent	excellent	excellent	good

Table 1: Summary of results for all experiments.

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 obtain an accurate distribution. 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.

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 likes. 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 (in contrast to complete stories) will be

important in realizing this goal.

In addition, we plan to incorporate this technique into the drama manager’s decision making process. 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 a principled trade-off between authorial intent and the player’s autonomy to pursue her own goals. Additionally, we intend to run a series of user studies to validate the assumptions we have made about player preferences and player satisfaction.

As noted, comparing the means of the NoDM and DM distributions does not provide an accurate measure of qualitative shift. At the top of our list for future work is to develop an meaningful summary measure that will enable us to make quantitative claims about the performance of our system. In addition to quantifying the accuracy of the estimation process, such a measure would also enable us to make quantitative claims about the effect of the drama manager.

Lastly, as mentioned above, we are very interested in taking a more in-depth look at how the set of features the player is trying to optimize relates to the quality of the approximation. Are there certain player features that are approximated well by the author’s features? Are there some that aren’t? Additionally, we also hope to better understand how the set of author features affects the quality of the approximation. Are there certain author features that approximate all player types well? Are there others that approximate specific players well? How does the number and diversity of author features affect the quality of approximation?

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. We also acknowledge the support of DARPA under contract No. HR0011-06-1-0021. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DHS, DOE, ORISE, or DARPA.

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/hcde.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)*.
- El-Nasr, M. S. 2007. Interaction, narrative, and drama creating an adaptive interactive narrative using performance arts theories. *Interaction Studies* 8(2).
- 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>.
- Kautz, H., and Allen, J. 1986. Generalized plan recognition. In *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86)*.
- 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*.
- Magerko, B. S. 2006. *Player Modeling in the Interactive Drama Architecture*. Ph.D. Dissertation, University of Michigan.
- 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)*.
- Picard, R. 2003. Computers that recognize and respond to user emotion. In *Proceedings of the International Conference on User Modeling*.
- Rich, E. 1979. User modeling via stereotypes. *Cognitive Sciences* 3(4):329–354.
- 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 trajectories in MDPs. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)*.
- Roberts, D. L.; Strong, C. R.; and Isbell, C. L. 2007. Estimating player satisfaction through the author's eyes. In *Artificial Intelligence for Interactive Digital Entertainment Workshop on Optimizing Player Satisfaction*.
- 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)*.
- Sweetsers, P., and Wyeth, P. 2005. Gameflow: A model for evaluating player enjoyment in games. *ACM Computers in Entertainment* 3(3):1–24.
- Thue, D., and Bulitko, V. 2006. Modelling goal-directed players in digital games. In *Proceedings of the second Artificial Intelligence and Interactive Digital Entertainment conference (AIIDE-06)*.
- Thue, D.; Bulitko, V.; Spetch, M.; and Wasylishen, E. 2007. Interactive storytelling: A player modelling approach. In *Proceedings of the third Artificial Intelligence and Interactive Digital Entertainment conference (AIIDE-07)*.
- Tsiriga, V., and Virvou, M. 2002. Initializing the student model using stereotypes and machine learning. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*.
- 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)*.