Emotion-Based Interactive Storytelling with Artificial Intelligence

Sergio Poo Hernandez and **Vadim Bulitko**
Dept. of Computing Science, University of Alberta
Edmonton, Alberta, T6G 2E8, Canada
pooherna@cs.ualberta.ca and bulitko@ualberta.ca

Emilie St. Hilaire
Dept. of Art and Design, University of Alberta
Edmonton, Alberta, T6G 2E8, Canada
esthilai@ualberta.ca

Abstract

Artificial Intelligence (AI) techniques have been widely used in video games to control non-playable characters. More recently, AI has been applied to automated story generation and game-mastering: managing the player's experience in an interactive narrative on-the-fly. Such methods allow the narrative to be generated dynamically, in response to the player's in-game actions. As a result, it is more difficult for the human game designers to ensure that each possible narrative trajectory will elicit desired emotional response from the player. We tackle this problem by computationally predicting the player's emotional response to a narrative segment. We use the predictions within an AI experience manager to shape the narrative dynamically during the game to keep the player on an author-supplied target emotional curve.

1 Introduction

Storytelling is a key art form. Stories are not only used to entertain but also to convey ideas, experiences and historical events with the intention of eliciting an emotion from the audience. Emotional response of the audience is a significant factor in managing the audience engagement and thus has implications for education, training and public policy decision-making. People expect an emotional connection from modern games to the point that a review of *Titanfall* (Respawn Entertainment 2014) criticized the expressly multiplayer game for lack of story-driven emotional engagement (Brown 2014).

The key matter to consider is the emotional response a particular person will have to a piece of narrative. Traditional writers employ several mechanisms (e.g., plot devices, character interaction) in an attempt to elicit a certain emotion from an *a priori* unknown reader. Filmmakers additionally use multiple emotional triggers such as background music, camera angles and colors to enforce the desired emotions (Tan and Fasting 1996; Smith 2003).

It is difficult enough to create a single static narrative that will influence a broad audience in a specific emotional way. Yet, modern narrative-rich video games bring in the added complexity of interactivity and hence emergent non-linear narratives. How can a video-game developer hope to elicit

certain emotions from an unknown player who is shaping the story dynamically by controlling their in-game avatar?

Video game developers have tackled this challenge in a variety of ways. One possibility is to make the narrative primarily emergent. Games such as the *Fallout* series (Bethesda Softworks LLC 2012) populate the game world with a great number of mini-narratives (side quests) and characters for the player to interact with in the hope that there is something for everybody in the game world. The downside of this approach is that the "main story" takes a backseat and usually lacks tension. For instance, in *Mass Effect 2* (Bioware 2010) the player is told that the galaxy is on the brink of destruction yet they solve domestic conflicts, unify broken families and enjoy dancing in a night club. An opposite approach is to write a linear story and hope that the topic is universal enough to elicit desired emotions from a broad range of players. A classic example is *Half Life 2* (Valve 2004) where the human kind is oppressed by technologically advanced aliens. Combining a narrative tension arc and open-world exploration is an outstanding challenge (Petit 2014).

Researchers in Artificial Intelligence (AI) have tackled these challenges by attempting to manage the player's experience on a per-player basis. Such AI experience managers approximate a human game master — a person who in the world of traditional pen-and-paper role-playing games assesses players' emotions and shapes the story to elicit certain emotional responses from them (Laws 2002).

If the story universe is encoded in a computer-readable form (e.g., first-order mathematical logic) then stories that are consistent with both the player's actions and authorial goals can be automatically generated via AI planners (Riedl et al. 2008). The downside of this approach is that it becomes more difficult for the game developers to anticipate and control the player's emotional response. To this end, AI researchers have attempted to model the player's state explicitly and shape the narrative specifically to influence it. For instance, PaSSAGE attempts to infer the player's inclinations towards different playstyles and then selects the next bit of narrative to maximize the player's perception of fun and agency (Thue et al. 2007; 2011).

In this paper we tackle the problem of automatically generating narrative on a per-player basis from a different angle. Answering the call for explicit emotion modeling in

AI-driven storytelling (Poo Hernandez and Bulitko 2013), we let the author explicitly specify a trajectory in a space of emotions that they would like their players to experience. So instead of providing concrete narrative goals (e.g., "Grandmother is eaten by the wolf" in an interactive version of "The Little Red Riding Hood" story (Riedl et al. 2008)) and hoping that accomplishing these will somehow elicit the needed emotional response from the player, we let the author specify emotional "key frames" at certain narrative points. For instance, the author can specify that at a certain point into the game the player should be hopeful but also somewhat afraid. Our AI system will then pick concrete narrative events which it expects will put the particular player in such an emotional state at the specified time.

We accomplish this goal by extending and combining several existing AI approaches to the interactive narrative experience management. Specifically, we encode the narrative world in a planning language and represent the stories themselves as plans in the spirit of the *Automated Story Director* (Riedl et al. 2008). We infer the player's inclinations towards different playstyles using the player model inspired by *PaSSAGE* (Thue et al. 2007). We use an appraisal model of emotions to predict the player's emotional response to a candidate narrative (Ortony, Clore, and Collins 1990; Marsella and Gratch 2003; Lazarus 1991). Our first contribution is thus a novel combination of several existing AI approaches and the introduction of a new computational model connecting playstyle modeling to goal inference to emotion prediction to narrative selection. Our second contribution is the specific implementation of the narrative planning module within our system. The resulting approach is called *Player Appraisal Controlling Emotions* (PACE).

We are in the process of implementing this approach in a unique testbed: possibly the first interactive narrative-oriented ballet video game. In addition to providing a rich testbed for PACE, the video game addresses such cross-disciplinary issues as exercise games (Lanningham-Foster et al. 2009) and gender representation and inclusion in video games (Norris 2004; Consalvo 2012).

The rest of the paper is organized as follows. We formulate the problem more precisely in Section 2. Section 3 reviews relevant existing work and its applicability to the problem at hand. Our approach is presented in Section 4. We discuss the on-going implementation within the interactive ballet video game and the planned empirical evaluation and the hypotheses it will test in Section 5. We then conclude the paper with directions for future work.

2   Problem Formulation

The problem we tackle in this paper is two-fold. First, we would like to give the player a sense of narrative agency by allowing them to meaningfully shape the story and get a feeling that their choices ultimately matter. Second, we would like the player to travel along an author-specified emotional trajectory. We attempt to solve both folds of the problem in the context of AI-managed interactive narrative where narrative events are produced as the game unfolds by an AI experience manager and are influenced by the player's choices,

expressed by his or her avatar's actions and the author's desires expressed in the form of a target emotional trajectory for the player.

In the terminology of the Procedural Game Adaptation (PGA) framework (Thue and Bulitko 2012), the video-game narrative is represented as a Markov Decision Process (MDP) with the player being the decision agent. The MDP states are narrative events and the MDP actions are the narrative choices the player makes. The MDP structure can be defined to ensure plot consistency (e.g., if the player kills the wolf early in the game, the little Red cannot be eaten by the said wolf later). Unbeknownst to the player, her actions are monitored by an AI experience manager which uses this information to adjust the MDP on the fly. This mapping from player's actions to the MDP adjustment is implemented using a variety of AI techniques as briefly outlined in the introduction and algorithmically detailed in the rest of the paper.

3   Related Work

Existing work relevant to the problem introduced in the previous two sections comes from two fields of research.

The first field includes AI experience management in the context of interactive narrative in video-game-like systems. The *Automated Story Director* (ASD) (Riedl et al. 2008) represents the narrative as a plan and uses an AI planner to build a narrative from a formal description of the story world and *a priori* given authorial goals. Not all plans (i.e., sequences of events) result in interesting and emotionally rich narratives. ASD lacks any model of the player's emotions or playstyle preferences. Instead, it forces the author to manually build an exemplar narrative. During the game, the player is monitored for rupturing the exemplar narrative by taking in-game actions. Such ruptures are then repaired by invoking an automated planner *Longbow* (Young 1994). While a narrative rupture can be repaired in many different ways, the planner selects the repair most similar to the handcrafted exemplar narrative thereby pulling the player back towards the original story. The lack of explicit player-specific emotion modeling prevents ASD from solving our problem by itself. We will, however, incorporate parts of ASD in our solution.

Explicit playstyle modeling in interactive storytelling was implemented in *Player-Specific Stories via Automatically Generated Event* (PaSSAGE) (Thue et al. 2007; 2011) where the player's actions were mapped to inclinations towards five distinct playstyles borrowed from pen-and-paper game mastering (Laws 2002). The inclinations were used to select from a handcrafted library of narrative segments in an attempt to maximize the player's feelings of enjoyment and agency. PaSSAGE neither allowed the game designer to specify an arbitrary emotional trajectory to keep the player on nor explicitly modelled the player's emotional state. While PaSSAGE cannot solve our problem by itself, we will use PaSSAGE-inspired playstyle model in our approach.

Player-specific Automated Storytelling (PAST) (Ramirez and Bulitko 2012; Ramirez, Bulitko, and Spetch 2013) combined the AI planner of ASD and the playstyle model of PaSSAGE in an attempt to repair player-induced ruptures in the exemplar narrative in a player-specific way. *Longbow* within PAST combined proximity to the exemplar nar-

rative with predictions of the player's enjoyment to select the best repair to a rupture. To predict the amount of fun the player would have along a candidate repair, PAST used a PaSSAGE-style model of playstyle inclinations, automatically updating it from the player's actions. This prediction of fun along a possible narrative can be viewed as rudimentary emotion modeling. However, it is insufficient to solve our problem as PAST did not allow the author to specify an arbitrarily shaped emotion curve even for fun, instead always attempting to maximize it. Furthermore, PAST could not be easily extended to support other emotions.

The other field of existing work focuses on inferring the player's emotional state (Lin, Spraragen, and Zyda 2012). For narrative-rich games the models based on goal-appraisal appeared a natural fit. Such models compute the player's emotional state as a result of the interaction between the player's goals and the likelihood of him or her achieving them. For example, the possibility of a failure elicits fear whereas a definitive failure elicits despair. A well-known appraisal model is OCC (Ortony, Clore, and Collins 1990). OCC is capable of modeling 22 different emotions and has been used in several systems such as EM (Reilly 1996), *Émile* (Gratch 2000) and *FearNot!* (Aylett et al. 2005; 2007). *Émile* computes the probability of an agent's success based on its current intentions and the plan the agent has developed to achieve those intentions and uses this probability to determine the agent's emotional state. *EMotion and Adaptation* (EMA) (Marsella and Gratch 2003) compliments an appraisal-based emotion modeling with a coping mechanism and thus can be used to control an NPC's appearance (Kenny et al. 2007) as well as actions within a game. While insufficient to solve our problem by itself, we incorporate an appraisal-style model into our solution.

A few projects attempted to shape narrative by explicitly representing the player's emotions. *Moe* (Weyhrauch and Bates 1997) used a target intensity curve and annotations on narrative events supplied by the author to guide the narrative. A similar approach is implemented in *Façade* (Mateas and Stern 2003) where each plot point is manually annotated with a tension value by the authors. Then an AI drama manager chooses the plot point whose tension value would be closest to the target tension curve. Our solution is similar but allows us to model a broader range of emotions and explicitly recognizes that the same plot point may elicit radically different emotional responses from different players. A different approach is used in *Distributed Drama Management* (DDM) (Weallans, Louchart, and Aylett 2012) where the non-playable characters model the player's current and future emotions and use them to choose an action to perform. Our solution is similar but is not character-centric.

With advances in biometric readers, researchers have attempted to explicitly read the player's emotional state and use it to shape the game. For instance, skin conductance, heart rate and facial electromyography (EMG) can be used to infer the player's level of tension and thus dynamically adjust the layout of a game level and enemies encountered by the player (Nogueira et al. 2013). A similar but sensor-less approach was implemented in the commercial game *Left4Dead* where the tension level of the player inferred

from observable variables (e.g., avatar's health level and shooting accuracy) was used to modulate the influx of enemies the player combats (Booth 2009). While the latter approach allowed the game designer to specify an arbitrary target tension curve, the changes to the game were limited to elements of the gameplay with no AI-managed narrative. An additional problem with biometrics-driven approaches is that only the player's current state can be read whereas generating future narrative requires a prediction of the player's future emotional response.

In summary, none of the related work reviewed in this section can solve our problem directly. However, most contain elements (e.g., playstyle modeling, narrative generation with AI planning) that can be incorporated into our solution.

4 Proposed Approach

As we described in the introduction, our solution, PACE, combines elements from existing AI experience managers in a novel approach. We will first illustrate the ideas with an example and then follow up with algorithmic details.
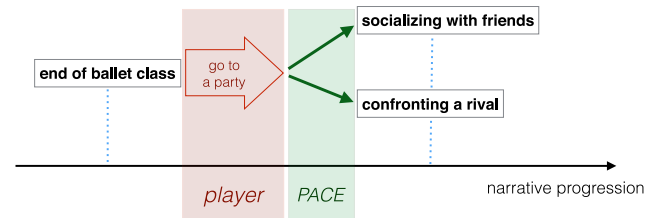
4.1 Intuition

Figure 1: The player goes to a party. What happens next?

Consider an interactive story based on the classic Romantic ballet *Giselle* (Gautier et al. 1841) where the player controls the titular heroine. At the end of a ballet class, the player decides to leave the studio (Figure 1) and head out to a party. It is now up to the AI experience manager to select the next bit of narrative for the player to experience. Using the automated planning approach from ASD, PACE computes two possible narratives with different scenarios upon arrival at the party: (i) the player socializes with her friends or (ii) confronts a rival dancer. Which one should be presented to the player?

As our goal is to elicit a specific emotional response from the player, the answer depends on the player's reaction to each of the two narrative candidates. To make the selection we predict the player's emotions using an appraisal model (Marsella and Gratch 2003). Suppose the author specified a target curve for the evolution of the player's hope, as shown in Figure 2. Then PACE will predict the emotional intensity of hope elicited in the player by socializing with friends as well as by confronting her rival and choose the narrative to keep the player closer to the target curve. This process is done in four steps as follows.

First, PACE maintains a model of the player's inclinations towards different playstyles (Thue et al. 2007). For this example, suppose there are three archetypical

playstyles: *storytelling*, *showing off*, *being modest*. Using author-supplied annotations on player's past actions, PACE has computed the player's inclinations towards these three playstyles as $(0.2, 0.1, 0.8)$.

Second, given the playstyle model of the player, PACE can now infer how desirable certain narrative goals are to her. Suppose the author had previously identified two possible goals that a player may be pursuing at this point of the story: *maintaining a successful career* and *leading a fulfilling personal life*. The author also provided a mapping between playstyle inclinations and goal desirability:

goal / inclination	storytelling	show off	modesty
career	2	3	0
personal life	2	0	3

Taking a matrix product of the mapping and the player's inclination

$$\begin{pmatrix} 2 & 3 & 0 \\ 2 & 0 & 3 \end{pmatrix} \times \begin{pmatrix} 0.2 \\ 0.1 \\ 0.8 \end{pmatrix} = \begin{pmatrix} 0.7 \\ 2.8 \end{pmatrix} \qquad (1)$$

PACE computes that the desirability of a successful career to the player is 0.7 whereas the desirability of a fulfilling personal life is 2.8.

Third, PACE uses the goal desirabilities and author-supplied probabilities of reaching these goals from each of the candidate narratives to predict player's emotional response. If the socialize-with-friends narrative is chosen then the probability of the player's having a successful career will be 0.3 while the probability of a fulfilling personal life will be 0.8. Using the appraisal model of emotions (Marsella and Gratch 2003), the emotion of hope is predicted to be elicited with an intensity of $0.3 \cdot 0.7 + 0.8 \cdot 2.8 = 2.45$.

On the other hand, the confronting-a-rival narrative places the player's chances of having a successful career and personal life at 0.9 and 0.2 respectively. Appraising the two goals gives the intensity of hope at $0.9 \cdot 0.7 + 0.2 \cdot 2.8 = 1.19$.
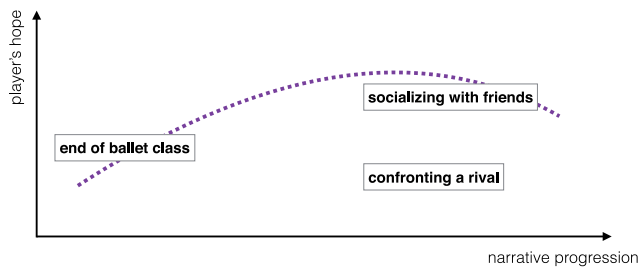

Figure 2: PACE selects "socializing with friends" narrative bit on the basis of keeping the player along a target emotional curve for the emotion of hope.

Fourth, PACE compares the predicted values of hope elicited by the two candidate narratives with the target value of hope that the author wanted the player to experience at that point of the narrative. If socializing with friends brings the player's hope closer to the target curve (Figure 2) then that bit of narrative will be presented to the player when they

arrive at the party. This process is applied to each scenario present in the narrative.

4.2 Algorithmic Details

In the MDP-based framework introduced in Section 2, PACE operates as shown in Algorithm 1.

Algorithm 1: PACE

inputs: narrative space (S, A, p), narrative start state s_1, narrative final states $S_f \subset S$, target emotion curve $\langle \bar{e}_t^* \rangle$

1. $t \leftarrow 1$
2. initialize playstyle inclinations $\bar{i}_1$
3. **while** $s_t \notin S_f$ **do**
4.   present narrative state s_t to the player
5.   collect player's narrative action a_t
6.   update playstyle inclinations $\bar{i}_{t+1}$ from a_t
7.   retrieve the relevant goal set G_t
8.   compute goal desirability $\bar{d}(G_t)$ from $\bar{i}_{t+1}$
9.   compute narrative candidates $\{n_j\}$ from s_t, a_t, p
10.   **for** *each n_j* **do**
11.   retreive goal probabilities $\Pr(G_t|n_j)$
12.   compute emotions $\bar{e}_j$ from $\Pr(G_t|n_j), \bar{d}(G_t)$
13.   compute deviation δ_j of $\bar{e}_j$ from $\bar{e}_{t+1}^*$
14.   select the smallest deviation: $j^* \leftarrow \arg\min_j \delta_j$
15.   select the next narrative state: $s_{t+1} \leftarrow n_{j^*}|_1$
16.   update the game dynamics p so that $s_t \xrightarrow{a_t} s_{t+1}$
17. $t \leftarrow t + 1$

As the inputs, the algorithm takes the set of narrative states S, the set of player's narrative-changing actions A, the MDP transition function p, the initial narrative state s_1, a set of final states $S_f \subset S$ and sequence of emotional states that the author would like the player to follow: $\bar{e}_t^*, t \in \{1, \dots\}$.

We first initialize the player's inclinations to different playstyles with a default inclination vector in line 2. The main loop of the algorithm (line 3) runs the story until the player arrives to one of the final narrative states: $s_t \in S_f$. Within the loop we present the player with the current narrative state s_t (e.g., "the ballet class ends") and collect his or her narrative input a_t (e.g., "go to a party").

In line 6 we update the player's inclinations towards different playstyles $\bar{i}_{t+1}$ as:

$$\bar{i}_{t+1} = \bar{i}_t + \bar{I}(a_t) \qquad (2)$$

where $\bar{I}(a_t)$ are author-supplied annotations. If any component of $\bar{i}_{t+1}$ exceeds 1 then we divide all of them by the maximum component.

We then retrieve the set of goals G_t the player may be pursuing at this point in the story (e.g., $G_t = \{$maintaining a successful career, leading a fulfilling personal life$\}$). In line 8 we compute the desirability $\bar{d}(G_t)$ of these goals using the updated inclinations $\bar{i}_{t+1}$:

$$\bar{d}(G_t) = \bar{d}(G) \times \bar{i}_{t+1} \qquad (3)$$

where $\bar{d}(G)$ is the author-supplied matrix linking playstyle inclinations to the goals as illustrated in Equation 1.

We then engage an automated planner (line 9) to compute possible narratives $\{n_j\}$ that follow the current narrative state s_t and the player's action a_t. Each narrative n_j is a sequence of possible future narrative states. In our example we had two narratives starting with the narrative states "socialize with friends" and "confront the rival" respectively.

Each of candidate narrative n_j is evaluated as follows. In line 11 PACE retrieves author-supplied probabilities of reaching the goals from the set G_t given the narrative n_j. For instance, the probability of reaching the goal "having a successful career" under the candidate narrative "socializing with friends" is 0.2.

We then use the appraisal model to compute the emotional intensities $\bar{e}_j$ (line 12) given the goal desirabilities $\bar{d}(G_t)$ and the probabilities $\Pr(G_t|n_j)$ of the narrative n_j allowing the player to achieve these goals. In line with the appraisal mechanisms of CEMA (Bulitko et al. 2008), our current implementation supports four emotions: distress, fear, hope and joy. Desirable goals elicit hope and joy. Undesirable goals elicit fear and distress and represented with a negative value. The intensity of an emotion is the product of the magnitude of the goal desirability/undesirability and the probability of reaching that goal. For instance, the intensity of hope elicited by the narrative n_j is calculated as the sum of the hope intensities for different goals from the set:

$$e_{\text{hope}} = \sum_{g \in G_t, \bar{d}(g) > 0, \Pr(g|n_j) < 1} \Pr(g|n_j)\bar{d}(g). \qquad (4)$$

Then the emotions predicted to be elicited by the narrative n_j form $\bar{e}_j = (e_{\text{distress}}, e_{\text{fear}}, e_{\text{hope}}, e_{\text{joy}})$. We can now compute the deviation of the predicted emotions $\bar{e}_j$ from the target emotional state $\bar{e}_{t+1}^*$ in line 13 as:

$$\delta_j = \|\bar{e}_j - \bar{e}_{t+1}^*\| \qquad (5)$$

where $\|\|$ is the 2-norm distance: $\|\bar{a} - \bar{b}\| = \sqrt{\sum_i (a_i - b_i)^2}$.

Once the candidate narrative that minimizes the deviation from the target emotion curve is selected (line 14), we set its *first* state as the next narrative state to be presented to the player (line 15). Following the non-stationary MDP framework, PACE modifies the MDP transition function p so that the player's action a_t leads her from the narrative state s_t to the first state of the chosen narrative: s_{t+1} (line 16). In the example above, the MDP transition function is modified by PACE so that when the player chooses to go to a party, she will socialize with her friends upon arriving there.

4.3 Implementational Details

While reusing elements from PAST, we replaced its automated planner, *Longbow*, with a domain-independent, PDDL-compatible planner. Switching to the *de facto* standard domain description language PDDL allowed us to take advantage of recent advances in automated planning research. We presently use FastDownward planner (Helmert 2006) running the LAMA (Richter and Westphal 2010) implementation. This planner was chosen because of its versatility and performance as it won the Sequential Optimization track and was the runner-up in the Sequential Satisficing track in the Seventh International Planning Competition

(IPC) (Coles et al. 2012). An example of encoding a domain action "dance" in PDDL is shown in Figure 3.

```
(:action dances
 :parameters (?p1 ?p2 ?p3 ?p4)
 :precondition (and (person ?p1) (person ?p2) (person ?p3)
                    (person ?p4) (not (knows ?p1 ?p2)))
 :effect (and (knows ?p1 ?p2) (emotion good) ))
```

Figure 3: A ballet action encoded in PDDL.

The planner is guided by a heuristic which PACE sets as the deviation δ_j of the emotions solicited by a candidate narrative n_j and the target emotional state. Thus, PACE is able to compute the best narrative n_{j*} without explicitly computing all alternatives $\{n_j\}$.

5   iGiselle

We are evaluating PACE in a novel testbed called *iGiselle*: an interactive version of the classic Romantic ballet *Giselle* (Gautier et al. 1841) set in a video game. An interactive ballet provides narratives with a variety of emotions, along with an opportunity to investigate gender issues in video games.

Figure 4: An *iGiselle* still image.

In *iGiselle* the player takes control of the titular heroine and experiences the narrative via a combination of still images (Figure 4), music and voiceovers. To further immerse
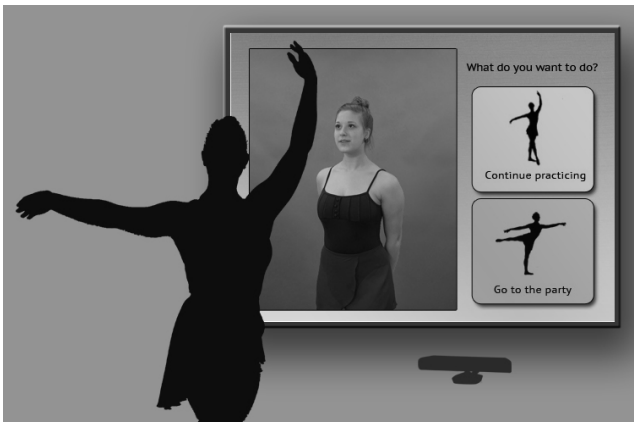
Figure 5: A mock-up of the *iGiselle* interface.

the player in the game we forgo a traditional game controller and have the player indicate their narrative choices by assuming dance positions (Figure 5) which are recognized with the Microsoft Kinect sensor.

The development of the multimedia content is done in two phases. First, working with writers we developed a non-linear narrative graph which allows the player to explore various narratives by controlling the heroine. The narrative graph is then encoded as states and actions in PDDL.

To capture the multimedia content for the narrative graph we have worked with ballet dancers, voice actors, choreographers and scriptwriters. The computer science part of our team is presently replacing the existing text-based interface of PACE with its multimedia counterpart. Once completed, the AI of PACE will drive *iGiselle* and generate narrative dynamically in an attempt to keep the player on an author-supplied emotional trajectory.

To evaluate the effectiveness of PACE we will conduct a user study in which human subjects will experience *iGiselle* managed by PACE using emotional modeling of the player (the experimental condition) or *iGiselle* managed by a baseline AI manager (the control condition). The subjects will then fill out a post-experience questionnaire on their enjoyment of the story, the emotions they experienced, etc.

6   Future Work

In addition to evaluating PACE with *iGiselle* as described above, we plan to explore applications of PACE to the so-called serious games. In particular, we are considering applications of emotion modeling to shaping training scenarios for neonatal emergency care. As with most emergency-response training, opportunities for live exercise are limited and mistakes can have catastrophic consequences (Bulitko and Wilkins 1999). As a result, medical instructors have shown interest in video-game-like intelligent training systems (de Ribaupierre et al. 2014). With PACE, an instructor will be able to specify the target emotional trajectory for the trainee, which PACE will then attempt to implement by dynamically shaping the training scenario in response to the trainee's actions.

Another interesting avenue for future work is development of narrative-rich exercise games. While a number of exercise-oriented video games have been developed (Morelli et al.), most of them do not immerse the player into a rich narrative. Using *iGiselle* we are planning to compare the level of story engagement the player would experience by controlling Giselle via our interface versus a gamepad.

Another area of improvement is how to obtain the goal desirability and goal likelihood values. Current implementation uses hand-coded values for the goal desirability and goal likelihood. But these values could also be obtained by interviews of test players and then relate their goal preferences to their playstyle inclinations. Another way is to set up test scenarios in which the players choices will be datamined for goal desirabilities. The evaluation of such methods would include their cost-effectiveness.

7   Conclusions

PACE is an AI experience manager that provides the player with a sense of narrative agency while keeping them on an author-specified emotional trajectory. In this paper we have shown how PACE is implemented in terms of the MDP-based framework PGA in order to accomplish these tasks. We also discussed an innovative narrative testbed in which PACE is being deployed.

8   Acknowledgments

References

Aylett, R. S.; Louchart, S.; Dias, J.; Paiva, A.; and Vala, M. 2005. FearNot!–an experiment in emergent narrative. In *Proceedings of Intelligent Virtual Agents*, 305–316.

Aylett, R.; Vala, M.; Sequeira, P.; and Paiva, A. 2007. FearNot!–an emergent narrative approach to virtual dramas for anti-bullying education. In *Virtual Storytelling. Using Virtual Reality Technologies for Storytelling*. 202–205.

Bethesda Softworks LLC. 2012. Fallout: New Vegas.

Bioware. 2010. Mass Effect 2.

Booth, M. 2009. The AI systems of Left4Dead. In *AIIDE*.

Brown, P. 2014. Titanfall Review - Through the Eyes of a Lapsed Combatant. http://www.gamespot.com/reviews/titanfall-review-through-the-eyes-of-a-lapsed-comb/1900-6415753/.

Bulitko, V. V., and Wilkins, D. C. 1999. Automated instructor assistant for ship damage control. In *IAAI*, 778–785.

Bulitko, V.; Solomon, S.; Gratch, J.; and van Lent, M. 2008. Modeling culturally and emotionally affected behavior. In *AIIDE*, 10–15.

Coles, A.; Coles, A.; Olaya, A. G.; Jiménez, S.; López, C. L.; Sanner, S.; and Yoon, S. 2012. A survey of the seventh international planning competition. *AI Mag.* 33(1):83–88.

Consalvo, M. 2012. Confronting toxic gamer culture: A challenge for feminist game studies scholars. *Ada: A Journal of Gender, New Media, and Technology* 1.

de Ribaupierre, S.; Kapralos, B.; Haji, F.; Stroulia, E.; Dubrowski, A.; and Eagleson, R. 2014. Healthcare Training Enhancement Through Virtual Reality and Serious Games. In *Virtual, Augmented Reality and Serious Games for Healthcare 1*. Springer. 9–27.

Gautier, T.; Coralli, J.; Perrot, J.; and Adam, A. 1841. Giselle.

Gratch, J. 2000. Émile: Marshalling passions in training and education. In *Proceedings of the fourth international conference on Autonomous agents*, 325–332.

Helmert, M. 2006. The Fast Downward Planning System. *J. Artif. Intell. Res.(JAIR)* 26:191–246.

Kenny, P.; Hartholt, A.; Gratch, J.; Swartout, W.; Traum, D.; Marsella, S.; and Piepol, D. 2007. Building interactive virtual humans for training environments. In *Proceedings of The Interservice/Industry Training, Simulation & Education Conference*, volume 2007.

Lanningham-Foster, L.; Foster, R. C.; McCrady, S. K.; Jensen, T. B.; Mitre, N.; and Levine, J. A. 2009. Activity-promoting video games and increased energy expenditure. *The Journal of pediatrics* 154(6):819–823.

Laws, R. D. 2002. *Robin's Laws of good game mastering*. Steve Jackson Games.

Lazarus, R. S. 1991. *Emotion and adaptation*. Oxford University Press New York.

Lin, J.; Spraragen, M.; and Zyda, M. 2012. Computational Models of Emotion and Cognition. In *Advances in Cognitive Systems 2*.

Marsella, S., and Gratch, J. 2003. Modeling coping behavior in virtual humans: don't worry, be happy. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, 313–320.

Mateas, M., and Stern, A. 2003. Façade: An experiment in building a fully-realized interactive drama. In *Proceedings of Game Developers Conference, Game Design track*, volume 2, 82–106.

Morelli, T.; Liebermann, L.; Foley, J.; and Folmer, E. An Exergame to Improve Balance in Children who are Blind.

Nogueira, P. A.; Rodrigues, R.; Oliveira, E.; and Nacke, L. E. 2013. Guided emotional state regulation: Understanding and shaping players affective experiences in digital games. In *AIIDE*.

Norris, K. O. 2004. Gender stereotypes, aggression, and computer games: an online survey of women. *Cyberpsychology & Behavior* 7(6):714–727.

Ortony, A.; Clore, G. L.; and Collins, A. 1990. *The cognitive structure of emotions*. Cambridge university press.

Petit, C. 2014. Balancing Freedom and Story in Dragon Age: Inquisition. http://www.gamespot.com/articles/ balancing-freedom-and-story-in-dragon-age-inquisition/ 1100-6419121/.

Poo Hernandez, S., and Bulitko, V. 2013. A Call for Emotion Modeling in Interactive Storytelling.

Ramirez, A., and Bulitko, V. 2012. Telling Interactive Player-Specific Stories and Planning for It: ASD + PaSSAGE = PAST. In *AIIDE*, 173–178.

Ramirez, A.; Bulitko, V.; and Spetch, M. 2013. Evaluating Planning-Based Experience Managers for Agency and Fun in Text-Based Interactive Narrative. In *AIIDE*.

Reilly, W. S. 1996. Believable Social and Emotional Agents. Technical report, Defense Technical Information Center (DTIC) Document.

Respawn Entertainment. 2014. Titanfall.

Richter, S., and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research* 39(1):127–177.

Riedl, M. O.; Stern, A.; Dini, D.; and Alderman, J. 2008. Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning* 4(2):23–42.

Smith, G. M. 2003. *Film structure and the emotion system*. Cambridge University Press.

Tan, E. S., and Fasting, B. T. 1996. *Emotion and the structure of narrative film: Film as an emotion machine*. Lawrence Erlbaum Associates, Inc.

Thue, D., and Bulitko, V. 2012. Procedural game adaptation: Framing experience management as changing an mdp. *AIIDE*.

Thue, D.; Bulitko, V.; Spetch, M.; and Wasylishen, E. 2007. Interactive storytelling: A player modelling approach. In *AIIDE*, 43–48.

Thue, D.; Bulitko, V.; Spetch, M.; and Romanuik, T. 2011. A Computational Model of Perceived Agency in Video Games. In *AIIDE*, 91–96.

Valve. 2004. Half-Life 2.

Weallans, A.; Louchart, S.; and Aylett, R. 2012. Distributed drama management: beyond double appraisal in emergent narrative. In *Interactive Storytelling*. Springer. 132–143.

Weyhrauch, P., and Bates, J. 1997. *Guiding interactive drama*. Carnegie Mellon University Pittsburgh, PA.

Young, M. R. 1994. A developer's guide to the Longbow discourse planning system. *University of Pittsburgh Intelligent Systems Program Technical Report* 94–4.