Proceedings of the Tenth Annual AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2014)

Exploring Abductive Event Binding for Opportunistic Storytelling

Emmett Tomai
University of Texas – Pan American, 1201 W. University Dr. Edinburg, TX, 78539
tomaie@utpa.edu

Abstract

We present a prototype game system using *opportunistic storytelling*, a particular commitment in the space of interactive narrative. It uses *abductive event binding* to deliver authored stories about the actions that the player is already taking to achieve game play goals. We show results of a simulation experiment that characterizes efficiency in delivering story event content, in terms of parameterized player motivation to follow storylines and the constraints placed on the events. Results show the value of two reasoning features in the system, *incentive* and *look-ahead*.

Introduction

Story is an important part of video games. Abstract rule systems are *framed* with setting, aesthetics, history and characterization to add meaning to the action. However, static narratives that work in framing (cf. Koster 2012) are unable to respond to the ongoing action that they motivate. This creates a frustrating disconnect between the agency players experience in game play and their passive audience role in the story (Jenkins 2004; Costikyan 2008).

Game designer Raph Koster defines the heart of game play as the exploration and mastery of the possibility space generated by a system of rules (Koster 2010). Good games have consistent rules that can be induced, game play that encourages practice and progress, and opportunities to exploit mastery in interesting ways. This fits the general definition of strong agency: awareness of alternative actions, perception of the connection from action to outcome, prediction of that connection and investment in the outcomes (Thompson et al. 1998). Game play has strong agency, but is repetitive and limited to what can be simulated. Narrative exposition can express the full range of human drama, but contradicts the player's agency.

There have been numerous attempts to give players agency over the direction of the story, from branching plots and conditional triggers to social simulations and AI research in human-like storytelling. However, they often fall into the trap of replacing game play with what we call *narrative choice*. Narrative choice occurs when player actions determine story direction based on authorial decisions rather than a consistent system of rules. Narrative choice is inherently weak agency, because the player cannot learn or predict what actions are available, or what their outcomes will be. The more the player tries to achieve specific outcomes, the more they will be forced to *play the narrative*, trying to guess what choices will get them where they want to go. Contrary to giving the player more agency over the story, playing the narrative is a poor game with little agency that also hurts immersion.

This problem is exacerbated when narrative choice has its own unique actions and decisions apart from game play. Consider stealing an item. As a narrative choice, there is no basis to make the decision apart from wondering where the story will go. But if stealing items is also part of game play, then there can be predictable, repeatable game play outcomes (e.g. being chased, losing "reputation"), and also unique, narrative outcomes (e.g. a bystander observes the crime and later confronts the player, advancing an authored plot). This fits well with a particular vision of interactive narrative, based on the premise that *stories often arise from unusual things happening to people going about their usual business*: the strange interaction on the train to work, the funny thing that happened at the store, and so on. Even fantasy heroes have the usual business of slaying monsters and rescuing people, intertwined with stories of a monster that isn't what it appears to be, or a victim caught up in a dark conspiracy. Telling the player unusual, unique stories about the usual, repetitive game play that they are already engaged in keeps decisions in context, mitigating the problems of narrative choice. We call this *opportunistic storytelling*. The guiding principles are:

1. All player actions result in consistent game play outcomes.

2. Game play goals are independent of narrative progress.
3. Narrative outcomes are presented when they fit with player actions.
4. Narrative outcomes do not interfere with game play outcomes.

From these principles, we present a prototype storytelling system that uses *abductive event binding* to find opportunities to *deliver* authored content that is consistent with player actions. We give the results of a simulation experiment showing how *incentive* and *look-ahead* interventions can make the system more efficient in delivering story content.

Related Work

Narrative choice works well if players are content with weak agency; exploring the designer's vision through their choices. TellTale Games' *The Walking Dead* is a well received narrative choice experience where players are encouraged to engage with the characters and setting, but make little difference in the story direction. In contrast, Quantic Dreams' *Heavy Rain* uses similar play mechanics, but player choices make significant changes in the story. *Heavy Rain* is a much more uneven experience, with players reporting feeling misled, not having "real" choices, or just playing the narrative to get the "right" ending. Narrative choice as a primary mechanic does not support strong agency.

Opportunistic storytelling builds on games such as EA/Maxis' *The Sims 2* and Lionhead Studio's *Fable* series, where game play actions trigger additional outcomes. *The Sims 2* has a library of short interaction patterns that it instantiates based on player actions to direct NPC agents. In *Fable*, player actions impact separate game systems for character morality and relationships. Both approaches were well received, suggesting promise for more complex, plot-oriented interventions.

AI research will continue to expand the boundaries of what can be simulated, turning narrative elements into game play elements. *Prom Week* (McCoy et al. 2013) makes social interaction into full-fledged game play with a set of underlying, consistent rules that the authors call *social physics*. There are clear objectives to adopt, and opportunities to explore, master and exploit the system. Other research seeks to formalize the process of generating dramatic narrative (cf. Szilas 2003), and where telling specific stories is not a priority, emergent narrative (cf. Louchart and Aylett 2004) is a viable alternative. But human-level interactive storytelling is more than simulation of reality and drama. A human storyteller can perceive the intentions behind player actions and generate outcomes that work with player expectations to give them real agency. AI has a long way to go to get to that point. *Façade* (Mateas and Stern 2003) is still arguably the most ambitious and complete playable interactive narrative experience from AI research. In it, the player can type free natural language dialogue at any time, and the system attempts to generate dramatic outcomes that are consistent with real life social interaction. *Façade* was well received as a novel experience and impressive system, but it cannot sustain coherent responses for the range of player input, much less work with player expectations to provide strong agency. In contrast, trainees in *IN-TALE* (Riedl and Stern 2006), a narrative choice military training exercise, have a strong context for their decisions – a well-defined process that they are highly motivated to follow and learn. Similarly, *Crystal Island – Outbreak* (Rowe et al. 2009) uses the *U-director* system (Mott and Lester 2006) to walk students through a research process to discover the cause of a mystery illness. That process provides context and external validation of player choices. In both cases, the experience is more even and less prone to play-the-narrative breakdowns. Opportunistic storytelling is an extension of this, using broad, general game play as that context. It is a pragmatic way to advance the field with engaging, playable experiences right now.

Opportunistic Storytelling

Evaluating the subjective quality of a game experience is a difficult task. However, before that point, there are quantifiable questions to resolve regarding the potential effectiveness of intelligent, non-guaranteed story delivery. In this work, we present simulation results to address the question of *efficiency*. Given that effort goes into creating stories, it matters what percentage of those stories will be seen by players. Stories here are abstracted to constrained *event patterns* with textual realization, as described below.

Consistent Game Play

Opportunistic storytelling is built on consistent game play, in this case a simple playable game in the style of a text-based Role-Playing Game. The game world is broken into a graph of *locations* with connecting *paths*. Each location has entities that afford the player certain actions. Figure 1 shows the simple interface used for development and testing. The main player activities are gathering items from *nodes* such as Berry Bushes, and fighting against *mobs* such as Bats. Combat is resolved through dice rolls, based on the fighters' health, which heals a small amount with each turn and fully when the player visits the Town area. Mobs can randomly attack the player when he or she attempts to take another action (e.g. gather). On defeat, they drop items for the player to collect. Both nodes and mobs disappear when used/killed and re-spawn after some time interval. In Town, the player is given goals to collect certain items. He or she then goes out into the wilderness to collect them, and returns to Town to make progress in the game. This is a deliberately simple system that gives the player consistent game play to pursue apart from story.

```
=========Forest=========
HP: 80
<Dark and scary>

The Forest's Edge is nearby.

There are 3 Herbs here.

There are 4 Bats here.

There are 2 Berry Bushes here.
===================
(e)xamine, (m)ove, (a)ttack, (g)ather,
(i)nventory, (q)uit:
```

Figure 1. Sample game interface

Story Delivery with Abductive Event Binding

While the player is going about his or her business, the game attempts to opportunistically deliver authored story content. Although it is up to the author, stories that are short and work together in a loose, environmental way are most likely to succeed. Every action the player performs results in one or more *events* being added to the game state. Those event outcomes are determined by the rules of the simulation. For example, a player attacking a Bear may result in an event where the player kills the Bear, or vice versa. If the player wins, the game displays the event with generic, repeatable text, e.g.:

```
You kill the Bear.
```

The storytelling system cannot alter this outcome, but can interpret the event, using abductive event binding.

Abduction is inference to the best, or in this case most convenient, explanation. For example, our system includes a rule that indicates that looting (taking items from a corpse) can also *produce* an examine event. That is to say, if you're pulling bits off a body, you might notice other things about it as well. This is not a deductive inference that says any time there is looting, there is also examining. Rather, if the system would like to observe an examine event, and there is a loot event, it is able to assume that an examine event did happen, subject to consistency checks. When a story event pattern is bound to an in-game event, it adds unique, non-generic narrative text, which goes outside the boundaries of the simulation, e.g.:

```
There's something very wrong with this Bear. It
looks like it's decaying from the inside out…
```

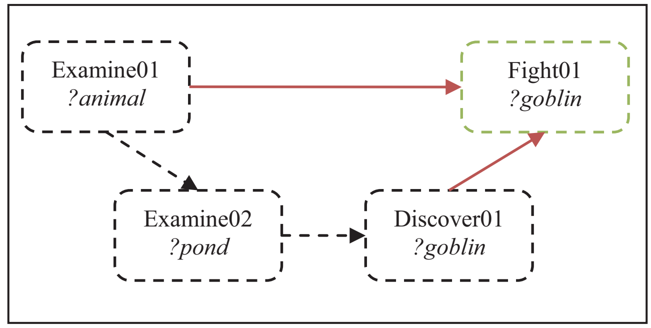

Figure 2. Event patterns and ordering for the Goblin Brewers story.

The story Goblin Brewers (GB) is depicted in Figure 2. The solid red arrows are required pre-requisites (disjunctive), while the dashed black arrows are ordering constraints. In this story, the player notices a sick animal (*Examine01*), finds a tainted pond (*Examine02*), discovers evil Goblins brewing a cauldron of poison nearby (*Discover01*) and defeats them (*Fight01*). Given the constraints, any of the first three events could be delivered first, but if *Examine01* or *Examine02* is skipped, then it is no longer available. These constraints are entirely up to the author, who can write a tight, detailed story with little flexibility, a loose story with many ordering options, or a story with many possible branches. The job of the system is to attempt to find opportunities to deliver as much of the story content as possible to the player within those constraints. The green border around *Fight01* indicates that it is marked as an *outcome* for this story. A story is incomplete until at least one outcome has been delivered. Figure 3 shows the first two event patterns in the GB story.

Examine01 (Type: EXAMINE)		
Slots	agent	PLAYER
	theme	?animal
	location	?loc1
Constraints	?loc1	WILDERNESS
Examine02 (Type: EXAMINE)		
Slots	agent	PLAYER
	theme	?pond
	location	?loc2
Constraints	?river	RIVER
	?pond	POND, TAINTED
	?loc2	WILDERNESS
	adjacentOrSame(?loc1, ?loc2)	

Figure 3. Event patterns from the Goblin Brewers story.

Given an in-game event, the system retrieves open (unbound) event patterns from stories in the story library. It retrieves them by event type (e.g. EXAMINE), ascending a hierarchy of specialization links in the knowledge base as necessary. The retrieve system also follows event productions. In our example story, if the player kills and loots a Bear, the system will retrieve *Examine01* and *Examine02*, because it can abduce an examine event from the loot event. Each retrieval creates a separate reasoning context where the system can try out bindings. The binding process follows these steps:

1) Ground entities in the in-game event are bound to the variables in the event pattern
2) Additional variables in the event pattern (e.g. *?pond* in *Examine02*) are bound to entities in that location
3) Constraints are checked

Binding fails if a variable is already bound to a different entity or an attribute constraint fails. For example, only entities with the POND attribute can bind to *?pond* in *Examine02*. The system can abduce attributes as necessary, provided they do not conflict with existing commitments. For example, an entity with the WOLF attribute cannot suddenly be given the POND attribute, which is enforced through the specialization hierarchy. Similarly, the attribute TAINTED is a VISIBLE_FEATURE, which cannot be assumed if the player can already see the entity. The variables in the event patterns are shared across the reasoning context at the story level. In the example, the *adjacentOrSame* constraint in *Examine02* means that the pond location *?loc2* must be adjacent to *?loc1* (or in the same location), which was bound in *Examine01* when the animal was found. But, if the system is binding *Examine02* without having *Examine01* (i.e. the player never found the animal) then *?loc1* is unbound and any pond anywhere will do. For competing binding opportunities, stories that are already in progress are favored over new stories, and they are scored by the number of steps left to reach an outcome.

In addition the event pattern outcome text, the author can add *hints* to the story that are tied to the ordering transitions. These are text snippets that alert the player to next possible steps in the story, which he or she can follow or ignore. For example, the transition between *Examine01* and *Examine02* has the hint:

```
Perhaps the sickness is something in the water.
```

By placing this hint on the transition, the author is indicating that the text guides the player to *Examine02*, and should be delivered when 1) the player has already experienced *Examine01* and 2) *Examine02* is possible in the current situation. To determine whether an open event pattern is possible, the system checks actions being presented to the user to see if they could bind. If so, the hint is presented.

Incentive and Look-Ahead

With no further intervention, the amount of story content that can be delivered is entirely dependent on what the player happens to do, and how much flexibility the author has allowed in the stories. Clearly, a story that allows an action to happen anywhere is more likely to be delivered than one that requires the player to stand in a certain location. But an intelligent human moderator would be able to increase the likelihood that opportunities for storytelling arise. They would not bind one event pattern in a location where the rest of the story isn't possible. And they would not send the player to an area where certain story elements cannot exist. We address this with two reasoning features: *incentive* and *look-ahead*.

Incentive

The incentive system takes control of random goal assignment to the player, and attempts to select goals most likely to result in storytelling opportunities. Prioritizing active stories first, it uses the same process that binds additional (non-slot) variables in event binding to find locations in the world where open patterns could be bound. For a given location, if the player knows what items are provided there, the system checks if any of those items are available as goals, and scores them as the inverse of the number of locations the player is aware of for obtaining that item. If the player has not been there, the system scores the items according to the actual number of locations where they are available. Those scores are combined with scores for the open patterns (inverse of how many prior open patterns would be skipped in the story) to decide which incentive goal to assign to the player. The incentive system should result in the player happening to be in the right place for the available stories more often.

Look-Ahead

The look-ahead system helps to decide whether it is a good idea to take an opportunity to deliver a piece of content when an event binding is available for it. The reasoning context for that proposed binding represents a world where the binding is accepted. Starting from that context, the system attempts to bind all the other open patterns in the story, scoring the initial binding according to how many subsequent bindings are possible. If the score of the proposed binding is less than a threshold, then the opportunity is passed on. In our example GB story, the system is less likely to invoke Examine01 if there is no location with a pond nearby. The threshold is set to the percentage of remaining goals, as a heuristic for how much longer the player will be playing in that area.

Experimental Setup

In this experiment, we evaluate the efficiency potential of this approach. For two specific stories, what is the likelihood that players will experience them completely or partially? The GB story has four events and strong topological constraints. The *Helpful Fairy (HF)* story, in contrast, is intended to be more flexible. It has only two events with limited constraints. Effectiveness in delivering a story is measured by whether the story is delivered *complete*, *abandoned*, *dead-end* or not at all. A complete delivery binds at least one event pattern that is marked as an outcome. An abandoned delivery is one where the player decided not to pursue the story further. When an incomplete story has at least one hint that was offered to the user, but not followed, it is considered abandoned. A dead-end delivery is incomplete but did not give the player any hints that were not followed.

There are many variables that impact the opportunities available to the system to deliver stories. In this experiment, the map layout, terrain features and story-relevant entities are fixed, while entities for game play are randomly placed. All adventuring takes place in the *Forest*, a four by four square of locations with the WILDERNESS attribute. There is one pond, one set of goblins in an adjacent location and one fairy. These locations are the same for every play through. Nine different types of gathering nodes and eight different types of mobs are distributed throughout the world before each play through. An experimental trial consists of three cycles of obtaining two goals, fulfilling them, and returning to the Town area outside the Forest. There are seventeen possible items to collect, and each can only be used once as a goal in a trial.

An automated player agent was created and used to run the experimental trials. The agent does not have access to the simulation data, and must explore the world to discover where items can be obtained. The agent is controlled by *motivation* parameters representing relative motivation for *achievement*, *exploration* and *story*. An achiever values completing goals efficiently by moving straight to known locations and obtaining items when possible. An explorer is motivated to visit unknown locations and examine entities of interest (the NOTABLE_FEATURE attribute stands in for callout text). A player with a higher exploration parameter is more likely to interact with the fairy, while a higher achiever is more likely to ignore the fairy to work on goals. The story motivation makes the player agent more likely to follow the hints provided by the system. To make a choice, the agent scores each available action provided by the simulation. Each action A_t is classified as one of eight decisions, shown in Table 1, and is scored according to:

```
if At == At-1:
  s = motiv(DCt) * (score(DCt) + stickiness(DCt))
else:
  s = motiv(DCt) * (score(DCt) + e)
```

Where A_{t-1} is the last action choice, DC_t is the decision class for A_t, *motiv* returns the agent's value for the motivation parameter associated with DC_t, and *score* and *stickiness* return the values from Table 1 for DC_t. Score is a relative priority for each class, and stickiness damps changing between decisions. The score and stickiness values are ad hoc, reflecting the simple game play choices.

Table 1.Action classifications for scoring.

Decision	Score	Stick-iness	Category
Examine point of interest	0.9	0.2	Explore
Explore Unknown Location	0.6	0.1	Explore
Follow story hint	1.0	0.2	Story
Gather	0.9	0.2	Achieve
Attack	0.85	0.2	Achieve
Turn In	0.8	0.5	Achieve
Move to Known Location	0.7	0.3	Achieve
Move to Known Location (before expected respawn)	0.5	0.1	Achieve

The experimental conditions are: base (*B*), with incentive enabled (*I*), with look-ahead enabled (*L*) and with both incentive and look-ahead enabled (*IL*). We ran 1000 trials of each condition, with random motivation parameters ranging from 0.5 to 1.5 for story and 0.8 to 1.2 for achievement and exploration to create variation. The narrower range for the latter was to reduce noise and focus on how motivation to follow story impacts delivery. We hypothesized that:

1) Higher story motivation leads to higher completion percentages
2) Less constrained stories (HF) are more likely to be completed than more constrained (GB)
3) The I and L conditions will improve completion percentages, and IL more so
4) More constrained stories (GB) will be impacted more by the incentive and look-ahead features

Results and Discussion

Table 2 shows the percentage of complete, dead-end and abandoned deliveries for the GB and HF stories. These results support hypothesis 2, showing that the HF story was completed 10-30% more often across the conditions. Hypothesis 3 is supported for GB, as the I and L conditions showed higher completion rates and the IL condition the highest. Hypothesis 3 was not supported for HF, as the L condition scored higher than B, but the I and IL conditions scored lower. Hypothesis 4 was somewhat supported as the I and IL conditions for GB clearly had more positive impact than for HF.

Table 2. Delivery type percentages for Goblin Brewers (GB) and Helpful Fairy (HF)

		Complete	Dead-end	Abandoned
GB	**B**	0.418	0.452	0.13
	I	0.471	0.391	0.138
	L	0.429	0.474	0.096
	IL	0.513	0.391	0.096
HF	**B**	0.737	0.05	0
	I	0.675	0.03	0
	L	0.754	0.047	0
	IL	0.676	0.027	0

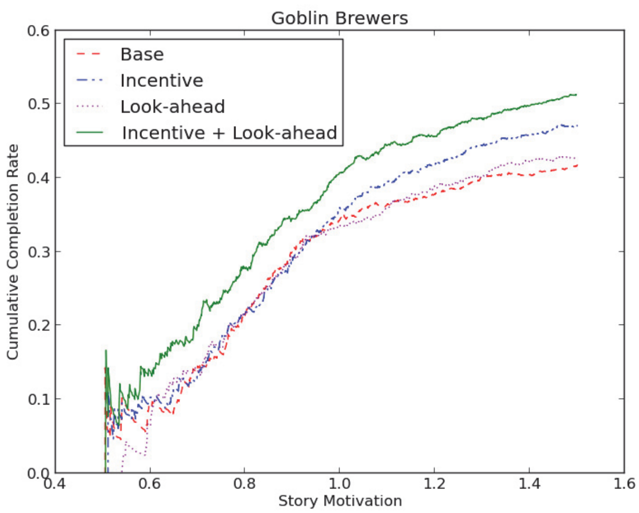

Figure 4. Cumulative completion rates for Goblin Brewers.

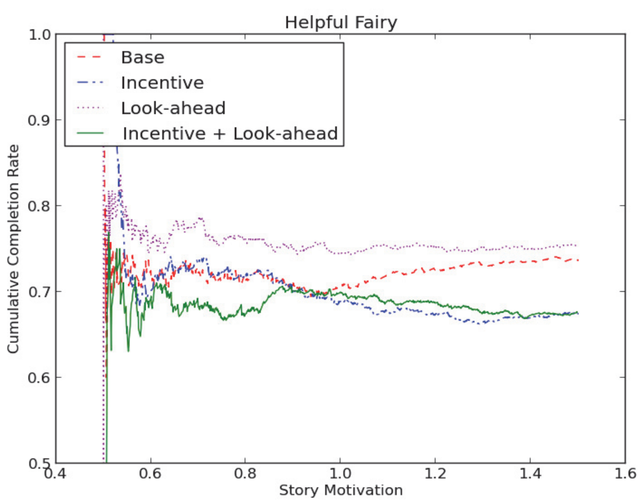

Figure 5. Cumulative completion rates for Helpful Fairy.

Figures 4 and 5 show the cumulative completion rates for GB and HF, accumulating the trials from lowest to highest agent story motivation. The GB case supports Hypothesis 1, as the completion rate improves as increasingly story-motivated trials are accumulated. It can also be seen that the IL condition outperforms the other conditions consistently across the spectrum of story motivation, even when the I and L conditions alone have no impact. Interestingly, the I condition gains advantage with higher story motivation. The HF case does not support Hypothesis 1, showing that an under-constrained story is just as likely to be completed or not, regardless of the agent's story motivation. Interestingly, the I and IL conditions get worse for more story-motivated agents.

Conclusion

From these results, we conclude that the system behaves reasonably with regard to simple player motivations and event constraints. This result holds up across randomly distributed game play goals (e.g. goals, nodes, mobs), suggesting that it is not dependent on players performing a specific set of actions or a specific level layout. It shows promise for working with different player motivation profiles, and different levels of story constraint. We also conclude that the incentive and look-ahead features are promising for building an intelligent story-telling system. By considering where the player is most likely to continue the story, the system is able to increase its efficiency in delivering complete stories. However, we did not predict, nor fully understand, the negative impact of incentive on the HF story, particularly when combined with look-ahead. Visualizing the experience of the player agents across large numbers of trials is a significant ongoing challenge.

Opportunistic storytelling brings together interesting areas of interactive narrative research with a well-defined experiential direction. First, there is the area of player prediction and goal recognition, which has been studied in the context of interactive narrative (cf. Magerko et al. 2004; Ha et al. 2011), and more general plan recognition (Charniak & Goldman 1993). Second, there is Thue's work on recognizing and using a model of player preference to inform content selection (Thue et al. 2007). Both directions are important for advancing this work, which currently does not attempt to recognize plan patters or player tendencies. Third, there is substantial research on character and relationship simulation to be leveraged to add additional playable narrative elements.

Other next steps for this work include moving to a game with more varied game play, enabling on-demand spawning of story-relevant entities, and building up a library of real stories to be told.

References

Charniak, E., & Goldman, R. P. (1993). A Bayesian model of plan recognition.*Artificial Intelligence*, *64*(1), 53-79.

Costikyan, G. 2004. Games, Storytelling, and Breaking the String. In N. Wardrip-Fruin and P. Harrigan (Eds.) *First Person: New Media as Story, Performance, and Game*, MIT Press.

Ha, E. Y., Rowe, J. P., Mott, B. W., & Lester, J. C. (2011, September). Goal Recognition with Markov Logic Networks for Player-Adaptive Games. In *AIIDE*.

Jenkins, H. 2004. Game Design as Narrative Architecture. In N. Wardrip-Fruin and P. Harrigan (Eds.) *First Person: New Media as Story, Performance, and Game*, MIT Press.

Koster, R. 2010. *A Theory of Fun for Game Design*. O'Reilly Media, Inc.

Louchart, S., & Aylett, R. 2004. Narrative theory and emergent interactive narrative. International Journal of Continuing Engineering Education and Life Long Learning, 14(6), 506-518.

Magerko, B., Laird, J., Assanie, M., Kerfoot, A., & Stokes, D. (2004). AI characters and directors for interactive computer games. *Ann Arbor*, *1001*(48), 109-2110.

Mateas, M. and Stern, A. 2003. Façade: An experiment in building a fully-realized interactive drama. Game Developers Conference, 2003.

McCoy, J., Treanor, M., Samuel, B., Reed, A., Mateas, M. & Wardrip-Fruin, Noah. "Prom Week: Designing past the game/story dilemma." *Proceedings of the 8th International Conference on Foundations of Digital Games*, Chania, Crete, Greece. May 14-17, 2013.

Mott, B. W., & Lester, J. C. 2006. U-director: a decision-theoretic narrative planning architecture for storytelling environments. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems* (pp. 977-984). ACM.

Riedl, M. O., & Stern, A. 2006. Believable agents and intelligent story adaptation for interactive storytelling. In *Technologies for Interactive Digital Storytelling and Entertainment* (pp. 1-12). Springer Berlin Heidelberg.

Rowe, J., Mott, B., McQuiggan, S., Robison, J., Lee, S., & Lester, J. 2009. Crystal island: A narrative-centered learning environment for eighth grade microbiology. In *Workshop on Intelligent Educational Games at the 14th International Conference on Artificial Intelligence in Education, Brighton, UK*(pp. 11-20).

Szilas, N. (2003, March). IDtension: a narrative engine for Interactive Drama. In*Proceedings of the Technologies for Interactive Digital Storytelling and Entertainment (TIDSE) Conference* (Vol. 3, pp. 187-203).

Thue, D., Bulitko, V., Spetch, M., & Wasylishen, E. 2007. Interactive Storytelling: A Player Modeling Approach. In *Proceedings of the Third Artificial Intelligence and Interactive Digital Entertainment Conference* (pp. 43-48).