Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence

Dramatis: A Computational Model of Suspense

Brian O'Neill
Western New England University
brian.oneill@wne.edu

Mark Riedl
Georgia Institute of Technology
riedl@cc.gatech.edu

Abstract

We introduce Dramatis, a computational model of suspense based on a reformulation of a psychological definition of the suspense phenomenon. In this reformulation, suspense is correlated with the audience's ability to generate a plan for the protagonist to avoid an impending negative outcome. Dramatis measures the suspense level by generating such a plan and determining its perceived likelihood of success. We report on three evaluations of Dramatis, including a comparison of Dramatis output to the suspense reported by human readers, as well as ablative tests of Dramatis components. In these studies, we found that Dramatis output corresponded to the suspense ratings given by human readers for stories in three separate domains.

Introduction

Narrative intelligence is an entity's ability to (a) organize and explain experiences in narrative terms, (b) comprehend and make inferences about narratives we are told, and (c) produce affective responses such as empathy to narratives. The pursuit of computational narrative intelligence has been a long-term goal of artificial intelligence. In this paper, we explore the question of whether a computational system can express narrative intelligence through the affective response to story content.

Suspense is an affective response, akin to anxiety, that humans frequently feel when being told a story. Expert storytellers who craft narratives for entertainment—films, novels, games, etc.—often structure their narratives to evoke strong affective responses. The idea that story structure is correlated with audience enjoyment dates back to Aristotelian notions of drama (Aristotle 1992) as well as more recent narrative theories (e.g., Freytag 1968). Just as expert human storytellers use suspense to maintain audience engagement, computational storytellers will need to do the same. As the use of virtual agents capable of telling stories grows, particularly in fields such as education and healthcare (Kenny et al. 2008), there will be a greater need for agents to engage and build a rapport with humans. Suspense is one of several storytelling techniques that such agents could use to keep stories interesting for their human audiences.

We present Dramatis, a computational model of suspense. Dramatis reads in a computational representation of a narrative structure and produces a suspense response that is comparable to suspense ratings given by human readers of an equivalent natural-language narrative. Dramatis uses a memory model to track a reader's conceptualization of the story as it is being read and uses scripts to anticipate undesirable outcomes for the protagonist. The model then searches for a plan to avert that negative outcome, where the perceived likelihood of success for that plan is correlated with the level of suspense.

Dramatis is a computational reformulation of a psychological model of suspense. While the psychological model provides a basis for Dramatis, intelligent systems require a degree of formalization that is typically unnecessary in a psychological model (Marsella and Gratch 2009). Thus, our reformulation more clearly defines the suspense phenomenon. We have evaluated Dramatis in three ways. We found that (a) Dramatis produces suspense responses that correspond to human self-reported suspense ratings on the same stories, (b) the memory model used by Dramatis is sufficient for producing human-like ratings, and (c) the methods used to define the planning problem are sufficient to match human suspense ratings.

Background and Related Work

Suspense is a technique used by authors and storytellers in a variety of narrative domains. There is no single unified definition of suspense, though definitions from narratology (Abbott 2008; Tan 1996; Branigan 1992), psychology (Gerrig and Bernardo 1994; Comisky and Bryant 1982; Ortony, Clore, and Collins 1988), and entertainment theory (Zillmann 1996) share common traits. Generally, these definitions claim that suspense requires uncertainty about a particular outcome (on the part of the audience), where the outcome is significantly desirable or undesirable.

This work is based on a modification of a psychological perspective of suspense, provided by Gerrig and Bernardo (1994). They argue that "readers feel suspense when led to believe that the quantity or quality of paths through the hero's problem space has become diminished." Gerrig and Bernardo describe a scene from Ian Fleming's *Casino Royale* novel, where James Bond, and readers thinking on his behalf, attempts to traverse a problem space from the current state (where a gun is held to his back), through a series of intermediate states, to a goal state (where his life is no longer in danger). Fleming introduces possible solutions, and immediately retracts them, thereby reducing the quantity of paths available to Bond (and the readers). Gerrig and Bernardo also note that authors can manipulate the search space by making certain actions appear too costly, thereby affecting the quality of paths in the space.

Of the various definitions of suspense, Gerrig and Bernardo's approach is the most amenable to a computational representation. The reader traversal of a problem space on behalf of a story character described by their model is analogous to the methods used in artificial intelligence to solve search problems. However, this model cannot be directly converted into a computational system; in the next section, we describe modifications made to Gerrig and Bernardo's model of suspense in order to represent it computationally.

Suspenser (Cheong 2007) is a discourse generation system that takes a given sequence of narrative events and produces the most suspenseful re-ordering of those events, possibly excluding events from the telling. Suspenser generates all possible plans a protagonist could use, and compares the ratio of failed plans to successful ones in order to evaluate the level of suspense. As the ratio of failed plans increases, the suspense level increases correspondingly. However, it is intractable to generate all possible plans for any sufficiently complicated story world model. Further, potential paths through the search space may terminate without reaching a solution for reasons that are specific to planning algorithms rather than the nature of the story world such as repeated states, lack of applicable operators, and problems binding literals to operators.

A number of story generation systems and interactive narrative systems use models of dramatic arc to guide the generation process. However, to date all of these systems require the human author to specify a fixed dramatic value for different events and actions that can be added to the story. These systems do not consider the question that events may appear more or less dramatic (or suspenseful) because of the context in which these actions or events are used. The interactive drama, Façade, tracks tension as part of managing the interactive narrative (Mateas 2002). The Façade drama manager tries to match the tension of each story state to an ideal tension curve. When the level of tension does not match the ideal curve, Façade probabilistically increases or decreases the tension in the story. The level of tension only affects which presentation of events a user sees, as each presentation conveys roughly the same events. Porteous et al. (2011) also model tension in their *Merchant of Venice* interactive storytelling system as a means of forcing stories to conform to traditional Aristotelian dramatic arcs. Interactive story authors can also use this system to design ideal dramatic arcs to serve as guides within their own narratives.

Reformulation of Suspense Definition

Creating a computational model of suspense requires a formal model of suspense. While Gerrig and Bernardo's model is suitable as a starting point, it is computationally intractable for three reasons. First, simply counting the paths that end in failures wrongfully assumes that all paths that do not lead to a solution are the result of the story state rather than the definition of the planning problem. Second, the definition suggests that humans regenerate the search space constantly and are in a perpetual state of feeling some level of suspense. However, most accepted definitions of suspense require that the audience expect some impending undesirable outcome in order to feel the phenomenon. Additionally, regenerating the search space requires consideration of the causal consequences of events, an inference which is typically possible only during offline processing (Graesser, Singer, and Trabasso 1994). Human readers are not capable of frequent offline processing, and therefore, a computational model of suspense should not regenerate the search space constantly. Finally, there is no evidence that human readers are capable of generating the entire search space, and thus all paths to success or failure, in part because humans are resource-bounded.

As a consequence of these insights into suspense and computation, we reinterpret Gerrig and Bernardo's description of suspense as a search space as follows:

> *Given the belief that the protagonist can face a negative outcome, one can assume that outcome will occur and search for the single most likely plan in which the protagonist avoids that outcome.*

We refer to such a plan as an *escape plan*. Further, where Gerrig and Bernardo refer to the quality of a path through the problem space, we define the quality of an escape plan as its perceived likelihood of success. Using perceived, rather than actual, likelihood allows us to consider the means available to authors for manipulating the search space. The perceived likelihood of an event is correlated with how quickly it can be retrieved from memory. Humans perceive the first thought they retrieve as being most likely to occur (MacLeod and Campbell 1992). Thus, a human reader would first generate the escape plan that they perceived as most likely to succeed. This approach does not require repeated generation of the search space, because the search is conducted only when the audience expects a negative outcome for the protagonist. Searching for a single plan avoids the need to generate the entire search space and overcomes the issues with measuring the ratio of failed plans.

Dramatis

Dramatis is a computational model of suspense that calculates the level of suspense over the course of a story according to the reformulation of Gerrig and Bernardo's definition of suspense presented above. Dramatis reads discretized symbolic-logic versions of stories, determines whether characters are facing a negative outcome, and generates and evaluates the quality of an escape plan to avoid that outcome. The evaluation of quality is equated to the level of suspense at that moment in the story.

Dramatis reads stories in a discretized symbolic-logic format called time-slices. Each time-slice describes one action in the story, where the action is given as an instantiated STRIPS operator. Time-slices also contain information about the characters in the scene, the location of the scene, and any effects of the action that cannot be inferred from the STRIPS operator.

Predicting Negative Outcomes

As Dramatis reads each time-slice, it predicts whether readers should expect any negative outcomes for the protagonist. To do this, the model relies on a library of script-like structures that represent the typical sequences in common story situations (e.g., a spy at a bar) and the negative outcomes that frequently occur within those situations (e.g., the spy getting poisoned). A script is a directed graph in which nodes are possible events and arcs are of two types. *Temporal arcs* indicate events that can follow other events. Any path through temporal arcs is a possible way the situation can unfold. *Causal arcs* indicate necessity relations between events in the same temporal arc path, and are labeled with the condition established by the first event that is necessary for the second event to happen. Dramatis uses the causal arcs to identify potential goals for the planning process in generating the escape plan.

When searching for relevant scripts, Dramatis prefers scripts that make use of recent events and can be bound to the actions in time-slices observed in the story so far. As the story continues, Dramatis tracks the identified script, maintaining a pointer to the most recently observed event that was also part of the script. In lieu of scripts, Dramatis can also glean similar information from opposing characters' plans that have been described within the story itself as part of the time-slices (e.g., a villain's monologue).

MEI-P Situation Model

As Dramatis reads the story, it adds story elements (such as characters and locations) and events (represented by STRIPS operators) to a memory model. We have implemented a memory model, *Modified Event Indexing with Prediction* (MEI-P), based on psychological theories of reader mental models. Zwaan et al. (1995) proposed Event Indexing (EI) to model the change in readers' conceptualizations of stories as they read. Niehaus (2009) proposed the Modified Event Indexing (MEI) model to account for narrative focus and the ability to draw inferences while reading. MEI-P uses scripts to activate concepts in memory based on future expectations.

The MEI-P situation model is a spreading activation network, where greater activation indicates that a story element is more salient in the reader's memory. Thus a story element with strong activation is more easily retrieved. The cost of retrieval will be used to calculate the perceived likelihood of actions used while constructing escape plans.

When Dramatis reads a new time-slice of the story, it creates a new node in MEI-P to represent the event in that time-slice. This event node is connected to nodes representing the characters in the scene, the location of the scene, the preconditions and effects of the event according to its STRIPS representation, and any literals used in the STRIPS representation. If such nodes did not already exist within the MEI-P situation model, then new nodes are added to the model. Each new edge in the network is given a weight of 1.0, while pre-existing edges decay according to their distance from the new event node. Node activation is calculated by giving all nodes an activation of 1.0, then iteratively spreading node weights according to edge weights until activation stabilizes.

In addition to incorporating story events observed in time-slices, MEI-P also includes events that are predicted by the scripts identified earlier in the process. Any event in the script that may occur after the most recently observed event is added to the MEI-P situation model. The salience of future events decreases with their distance from the current event, just as prior events become less salient over time unless connected to more recent events.

Generating Escape Plans

Dramatis' suspense response is correlated with the perceived likelihood of success for escaping an impending negative outcome. Dramatis generates the escape plan—a plan that averts the negative outcome—that is perceived to be most likely to succeed, and the cost of the plan is equivalent to the level of suspense. Dramatis generates an escape plan as follows. First, the currently activated script is converted into a planning problem. The initial state is the current state of the story world as determined by the logical formalism of the story. The goal situation is the negation of any condition on any causal arc in any valid path through the active script between the current state and the undesirable outcome. That is, any sequence of actions by the hero that violates a causally necessary condition leading to a negative outcome may avert that negative outcome.

The planning problem is solved with a modified version of the Heuristic Search Planner (HSP) (Bonet and Geffner 2001), which can find near-optimal solutions for problems with non-uniform costs. HSP is an informed state-space search algorithm that uses a relaxed form of the planning problem—ignoring delete lists in the STRIPS operators—to estimate the cost of reaching a goal state. HSP requires that the goal be stated as positive (non-negated) propositions. Because our goal situation is comprised of a number of negated propositions ORed together, we modify HSP to allow it to peek at the delete lists to determine if any operator can achieve the goal situation. HSP is not guaranteed to return the optimal plan in our case, but typically converges on a relatively low-cost solution.

Operator cost is calculated using the activation levels of corresponding nodes in the MEI-P situation model. Because Dramatis uses perceived likelihood as a measure of plan quality, we assume that the easiest thing for a reader to recall from memory is perceived as most likely to succeed (MacLeod and Campbell 1992). That is, elements that are strongly activated in MEI-P are easier to retrieve. Therefore, an operator's cost is inversely correlated with the activation of the elements within the operator, including the operator itself, the propositions representing the operator preconditions and effects, and the literals used within the operator. The cost of a plan is equivalent to the sum of the costs of each action within the plan.

Dramatis tracks escape plans that were generated as it continues to read the story. When the events in subsequent time-slices match the events predicted in the escape plan, Dramatis recalculates the cost of the remainder of the escape plan rather than generating a new one. When newly observed events do not overlap with the escape plan, Dramatis starts the escape plan search process over again. While this may result in the same plan, the cost of the plan, and therefore the level of suspense, will likely differ.

After each time-slice, the cost of the resulting escape plan is equivalent to the level of suspense. Over the course of a story, Dramatis creates a *suspense curve*, showing the change in suspense over time.

Evaluations

We conducted three evaluations of the Dramatis model. In the first evaluation, Dramatis generates suspense curves for two versions of three stories, while human readers read and rate the suspense levels of natural language versions of these same stories. We report the level of agreement of suspense ratings between Dramatis and human readers. The second evaluation ablates the MEI-P situation model, such that it is not used in calculating escape plan cost, and therefore, the suspense rating. The suspense curves generated by the ablated system are compared to the human ratings generated in the first evaluation. In the final evaluation, we modify the algorithm for determining goal states prior to escape planning, in order to demonstrate the necessity of the goal selection process.

These evaluations required pairs of stories with differing levels of suspense. Thus, the evaluations use two versions of three stories. Each pair contains an original version and an alternate version, where details of the original version were changed in a way that we believed would lower the overall perception of suspense. The three story pairs were adapted from scenes from the films *Casino Royale*, *Rear Window*, and *Harry Potter and the Half-Blood Prince*.

The alternate versions of the stories were created by removing, reordering, or changing the details of the events of the original version. The changes reduce suspense by providing additional means of addressing the protagonist's problem or introducing solutions earlier. For example, in the original *Casino Royale* scene, James Bond is poisoned and attempts to cure himself. In the alternate version, we introduce the antidote sooner, and we remove Bond's failed attempt at curing himself. As a result, readers can use the antidote in solutions sooner and retain multiple possible solutions.

The operators and scripts used by Dramatis were crowdsourced to avoid experimenter bias; details of our crowdsourcing methodology for crowdsourcing scripts and operators are beyond the scope of this paper. The operators, scripts, and all versions of the stories used in these evaluations are available in (O'Neill 2013).

Evaluation 1: Comparison with Human Readers

In Evaluation 1, Dramatis generates suspense curves for both versions of the three stories. Human readers read natural language versions of each pair of stories, rating the suspense of each version, and selecting which member of a given pair was more suspenseful. We hypothesized that Dramatis and human readers would produce the same within-pair ordering of stories according to suspense level.

Method

Thirty-two participants were recruited to read natural language versions of the stories described above. All participants read the story pairs in the same order (*Casino Royale, Rear Window, Harry Potter*), but we controlled for reading order within-pairs (original vs. alternate versions).

Table 1. *Suspense ratings as reported by readers and Dramatis in Evaluations 1-2. Asterisks indicate statistical significance.*

	Reader Preference		Reader Rating (Interpolated Median)		Dramatis Area Under Curve		Dramatis (Without MEI-P) Area Under Curve	
	Original	Alternate	Original	Alternate	Original	Alternate	Original	Alternate
Casino Royale	31*	1	5.59*	3.23	20526	7468	35.5	24
Rear Window	18	14	4.90*	4.73	4176	1170	53.5	53.3
Harry Potter	23*	9	4.83*	4.14	27428	16293	64	50

After reading each individual story, participants answered the question, "How suspenseful was this story?" on a 7-point ordinal Likert scale. After each pair of stories, participants indicated which version of the story was more suspenseful, thus producing an ordering for the pair of stories according to suspensefulness. Participants were not able to see previous responses while answering questions.

Dramatis was given the same six stories in Time-Slice format. Because Dramatis produces suspense curves, rather than overall ratings, we measure the overall suspense of a particular story according to the area under the curve.

Results and Discussion

Table 1 shows the results of Evaluation 1. For each of the three pairs of stories, the area under the Dramatis suspense curve for the original version was greater than the area under the alternate version curve.

When asked to choose which of the two story versions was more suspenseful, human readers selected the original version for all three pairs of stories. The preference rates for the original version were statistically significant for *Casino Royale* ($p < 0.001$) and *Harry Potter* ($p < 0.05$). When providing ratings for each story, participants provided higher suspense ratings for the original version of each pair. The differences in ratings were statistically significant for *Casino Royale* ($p < 0.001$), *Rear Window* ($p < 0.05$), and *Harry Potter* ($p < 0.01$).

These results confirm our hypothesis. Human readers consistently rated the original version to be more suspenseful than the alternate version of each pair. Furthermore, Dramatis generated suspense curves with greater areas underneath the curve for the original version of the story.

Evaluation 2: Ablation of MEI-P

In Evaluation 2, the MEI-P situation model is removed from the calculation of planning operator costs. Instead, all planning operators have equal cost during the escape plan process. Thus, without MEI-P, suspense ratings produced by Dramatis correspond to the *length* of the escape plan. We hypothesize that this modified version of Dramatis will produce an ordering of stories according to suspensefulness that is inconsistent with the ordering provided by human readers in Evaluation 1. If this hypothesis holds, we can conclude that the MEI-P Situation Model is a sufficient component for Dramatis to produce suspense ratings that agree with those provided by human readers.

Method

The modified version of Dramatis was given the same three pairs of stories as in Evaluation 1. All other input was identical. Suspense ratings for the stories are again calculated using area under the curve.

Results and Discussion

The last two columns of Table 1 show the areas under the curve for the suspense curves produced by the modified Dramatis system. In the case of *Casino Royale*, the area under the curve was greater for the original version than for the alternate version. However, the curves produced are largely identical, with the exception of a three time-slice section where the original version is considered more suspenseful. For *Rear Window*, the ablated Dramatis system generated identical suspense curves for the original and alternate versions. Thus, the two versions are considered equally suspenseful. Finally, the modified system produces suspense curves for *Harry Potter* with a larger area under the curve for the original version.

Our hypothesis is partially supported. The modified Dramatis produced identical suspense curves for the two versions of *Rear Window*, failing to match the ordering produced by human readers. However, despite not having the MEI-P situation model, the modified Dramatis produced the same ordering for *Casino Royale* and *Harry Potter* as human readers. The differences between the two *Casino Royale* versions were entirely isolated to a small section of the story, whereas for the most part, the model found the two versions to be largely identical. Given that the modified model could not entirely replicate the ratings produced by human readers, we believe that the MEI-P situation model is a sufficient component for producing suspense ratings that correspond to those of human readers.

Evaluation 3: Modification of Goal Selection

In Evaluation 3, we modify the procedures used to select goal states during the escape planning process. Rather than selecting the causal link that leads to the most easily retrieved escape plan, the modified Dramatis model randomly selects a link from the set of candidate causal links to be the goal situation for the escape planning problem. By altering this portion of the Dramatis algorithm, we can demonstrate that the original goal selection procedures lead to consistently correct orderings of story suspense. Conversely, the modified strategy leads to incorrect and inconsistent interpretations of the suspensefulness of stories. We hypothesize that this modified form of Dramatis is unable to consistently match the orderings of stories according to suspense level produced by human readers in Evaluation 1.

Method

The modified version of Dramatis was given the same three pairs of stories as in Evaluation 1. All other input was identical. Suspense ratings for the stories are again calculated using area under the curve.

The goal selection algorithm was modified as follows: At each time-slice, Dramatis randomly selects one of the candidate causal links, generates an escape plan for that particular goal state, and calculates a suspense rating. This rating is added to the suspense curve. For each story version, we sampled 1000 curves generated using this process and compared curves across story versions using area under the curve. Additionally, we calculate the likelihood that the alternate version of the story is rated more suspenseful than the original version at particular points of each story.

Results and Discussion

When comparing sampled suspense curves from the two versions of *Casino Royale*, the curve for the original version had a larger area than the alternate version of the curve in 92.3% of cases ($p < 0.001$). For *Rear Window*, the area under the original version curve was greater 99.8% of the time ($p < 0.001$). For *Harry Potter*, the area under the original version curve was greater 62.9% of the time ($p < 0.001$). Thus, for each pair of stories, randomly generated suspense curves tended to indicate that the original version of a story was more suspenseful than the alternate version.

When looking at specific points of each story, we can see that there are problems with the random selection process. For example, with each pair of stories, we observed examples where one version of a story would be rated more suspenseful than the other, at times when we would expect the two stories to be equally suspenseful (e.g., near the beginning of the stories, before the story pairs diverge). At other points, we observed situations where the alternate version of the story was rated more suspenseful, even though the baseline version of Dramatis indicated that the original version was much more suspenseful.

The results of sampling curves indicates that, by taking the “majority vote” of a large sample of curves, we can produce the same suspense rankings as human readers. A randomly generated pair of curves would select an ordering that matched the human ordering between 62%-99% of the time, depending on the story domain. Therefore, the algorithm for deterministically selecting causal links to negate is not the only algorithm that is capable of matching human suspense rankings. However, note that the random sampling strategy overwhelmingly selected the original version of *Rear Window*, while human readers struggled to reach consensus about which version was more suspenseful. This may indicate weaknesses in the sampling strategy. Additionally, this random selection strategy leads to flawed situations, such as those noted above, where one version of a story is deemed more suspenseful when the two stories should be considered equally suspenseful. Our hypothesis for Evaluation 3 is considered partially supported, as the modified system did not consistently select the original version of each pair to be more suspenseful, although a sampling strategy was devised that did match the ordering given by human readers.

These evaluations demonstrate that the suspense ratings produced by Dramatis are largely consistent with those produced by human readers. Further, when provided with stories intended to be less suspenseful, Dramatis produced correspondingly lower ratings. The latter evaluations also demonstrated that without the MEI-P situation model or the goal selection procedures for escape planning, Dramatis is not capable of consistently producing suspense ratings corresponding to those produced by human readers.

Conclusions

The Dramatis model is based on a reformulation of a psychological model of suspense. Under the original model, reader suspense is correlated with the appearance of diminishing quantity or quality of paths through a hero’s problem space. In our reformulation, the level of suspense is correlated with the perceived likelihood of success of a plan generated to get the protagonist out of a dire situation. Our evaluations of Dramatis demonstrated that the model and human readers produce the same ordering of stories according to suspense level. Additionally, these evaluations showed the sufficiency of the reader situation model and the processes used for generating escape plans.

Dramatis serves as an initial step towards a computational model of aesthetics. While suspense is hardly the only consideration when evaluating the aesthetics of stories, this model represents progress towards a broader model of human aesthetics. This particular model could be applied to computationally creative systems, in fields such as story generation (Gervás 2009) or interactive narrative (Riedl and Bulitko 2013). Such systems could iteratively use Dramatis to evaluate the suspense level of stories, and change the story to reach a target level of suspense, continuing until some threshold of suspense had been achieved.

Dramatis demonstrates narrative intelligence, producing affective responses (in the form of suspense) in response to story content. This model can serve as a guide for future efforts to formalize psychological models of narrative intelligence in computational representations. With continued pursuit, it may be possible to generate and respond to stories with human-like affective responses.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 1002748. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Abbott, H. P. 2008. *The Cambridge Introduction to Narrative*. Cambridge, UK: Cambridge University Press.

Aristotle. 1992. *The Poetics*. (T. Buckley, trans.) Buffalo, NY: Prometheus Books.

Bonet, B., and Geffner, H. 2001. Planning as Heuristic Search. *Artificial Intelligence* 129:5-33.

Branigan, E. 1992. *Narrative Comprehension and Film*. New York: Routledge.

Cheong, Y. 2007. A Computational Model of Narrative Generation for Suspense. Ph.D. diss., Department of Computer Science, North Carolina State University, Raleigh, NC.

Comisky, P., and Bryant, J. 1982. Factors Involved in Generating Suspense. *Human Communication Research* 9:49-58.

Freytag, G. 1968. *The Technique of Drama: An Exposition of Dramatic Composition and Art*. Johnson Reprint Corporation.

Gerrig, R. J., and Bernardo, A. B. 1994. Readers as Problem-Solvers in the Experience of Suspense. *Poetics* 22:459-472.

Gervás, P. 2009. Computational Approaches to Storytelling and Creativity. *AI Magazine* 30(3): 49-62.

Graesser, A.C.; Singer, M.; and Trabasso, T. 1994. Constructing Inferences During Narrative Text Comprehension. *Psychological Review* 101:371-395.

Kenny, P.; Parsons, T.; Gratch, J.; and Rizzo, A. 2008. Virtual Humans for Assisted Health Care. In *Proceedings of the 1st International Conference on Pervasive Technologies Related to Assistive Environments*. New York: ACM Press.

MacLeod, C., and Campbell, L. 1992. Memory Accessibility and Probability Judgments: An Experimental Evaluation of the Availability Heuristic. *Journal of Personality and Social Psychology* 63:890-902.

Marsella, S. C., and Gratch, J. 2009. EMA: A Process Model of Appraisal Dynamics. *Cognitive Systems Research* 10:70-90.

Mateas, M. 2002. Interactive Drama, Art and Artificial Intelligence. Ph.D. diss., Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA.

Niehaus, J. 2009. Cognitive Models of Discourse Comprehension for Narrative Generation. Ph.D. diss, Department of Computer Science, North Carolina State University, Raleigh, NC.

O'Neill, B. 2013. A Computational Model of Suspense for the Augmentation of Intelligent Story Generation. Ph.D. diss., School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA.

Ortony, A.; Clore, G. L.; and Collins, A. 1988. *The Cognitive Structure of Emotions*. Cambridge, UK: Cambridge University Press.

Porteous, J.; Teutenberg, J.; Pizzi, D.; and Cavazza, M. 2011. Visual Programming of Plan Dynamics Using Constraints and Landmarks. In *Proceedings of the 21st International Conference on Interactive Digital Storytelling*, 186-193. Menlo Park, CA: AAAI Press.

Riedl, M. O., and Bulitko, V. 2013. Interactive Narrative: An Intelligent Systems Approach. *AI Magazine* 34(1): 67-77.

Tan, E. S. 1996. *Emotion and the Structure of Narrative Film: Film as an Emotion Machine*. New York: Routledge.

Zillmann, D. 1996. The Psychology of Suspense in Dramatic Exposition. In *Suspense: Conceptualizations, Theoretical Analyses, and Empirical Explorations*, 199-231. Mahwah, NJ: Lawrence Erlbaum Associates.

Zwaan, R. A.; Langston, M. C.; and Graesser, A.C. 1995. The Construction of Situation Models in Narrative Comprehension: An Event-Indexing Model. *Psychological Science* 6:292-297.