A Computational Model of Emotional Response to Stories

Adam Fitzgerald, Gurlal Kahlon, and Mark O. Riedl

School of Interactive Computing, Georgia Institute of Technology
Atlanta, Georgia, USA
{afitz, kahlon.gurlal, riedl}@gatech.edu

Abstract. In this paper, we consider the problem of computing the affective responses that humans experience when reading or watching stories. Evidence suggests that emotional responses result from recipient (reader, watcher, etc.) problem solving on behalf of story world characters when the recipient predicts undesirable narrative outcomes. Our system computes the level of tension a human recipient is expected to experience as a narrative unfolds. The system efficiently determines possible future outcomes, measures their utility, and estimates the probability that they will occur. The resultant estimate of tension is a function of expected utility.

1 Introduction

Humans actively reason about stories [1], and responses can be categorized as cognitive (e.g, belief change) and affective (e.g, emotional response, tension, etc.). This later aspect, affective response, has received less attention from the research community. In this paper, we present work toward a model of affective response to narrative. Building off of narratological and psychological theories of suspense [1–4], we formulate affective response as a function of expected utility of predicted future narrative outcomes. Narrative recipients (readers, watchers, listeners, etc.) actively reason on behalf of the protagonist [1]. Suspense is an emotional response that manifests itself physiologically as a feeling of anxiety or tension. One cognitive account of suspense posits that it occurs when a narrative recipient believes that the likelihood of avoiding a negative future state is perceived to be small or nonexistent [2]. As options for avoiding the negative future state are removed, and the probability of the negative future state increase, feelings of tension increase correspondingly. Our approach correlates tension to expected utility. Our *Response Model* forecasts future states, determines their value, and estimates the probability that undesirable future states will be averted by character actions or happenstance. These computations yield a result that can be interpreted as relative tension.

2 The Response Model

To compute the degree of tension that a human recipient of narrative is expected to experience, the system must perform the following:
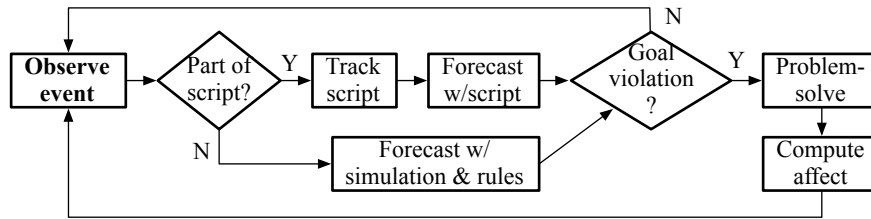
Fig. 1. Diagram of the Reader Model process.

1. **Adopt goals.** The model determines which characters the recipient will have have high affinity for. It adopts the goals of the character (including *maintenance goals* such as "not be harmed").
2. **Forecast future states.** The model forecasts possible future states, looking for violations of adopted character goals. The utility of these future states is determined.
3. **Estimate the probability of future states and compute tension.** Should an adopted goal be violated, the system estimates the probability that the predicted future state will be avoided. The probability estimation is combined with utility to determine relative level of tension.

Fig. 1 shows how the Response Model processes relate. The story is fed into the Response Model in the form of text structured as *frames*. Each frame of the story provides (a) characters and their characteristics; (b) the action of the character; (c) the location; and (d) any extra-deigetic information such as background music. Frames are structured chunks of narrative designed to correspond to the information that a human recipient can extract from a sentence (or a couple of sentences together representing an event) or a beat in a film. Affinity for characters is calculated dynamically as a function of the valence of character actions and the number of references to the character in the narrative, modulated over time to take on an S-curve shape with an asymptotic upper bound. The Response Model adopts any goals associated with characters with high affinity.

2.1 Predicting the Future

The Response Model uses three cognitively plausible mechanisms by which to predict future world states: scripts, simulation, and rules. Scripts are schema-like representations of familiar situations used to reduce the cognitive load of acting in the world [5]. For simplicity, scripts are represented as finite state machines. The FSM representation deterministically captures the variations that can occur in a script. Each script has exactly one entry state – called the triggering state – and exactly one exit state. A script is loaded and tracked if any situation matches the triggering state. The exit state of the FSM is annotated with state information for involved characters that can be compared to adopted goals. Fig.
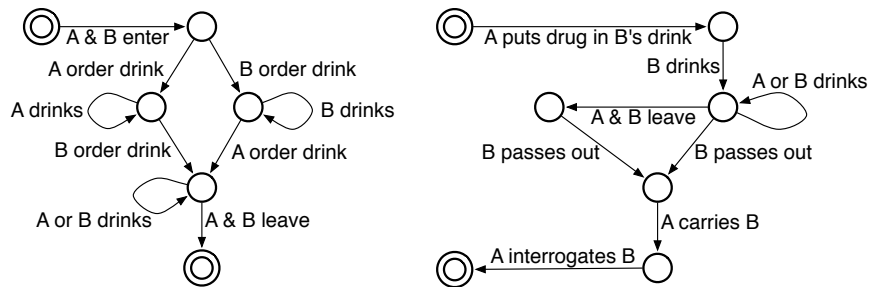
Fig. 2. Finite state machine representation of scripts representing (a) two people having drinks at a bar, and (b) a kidnapping via drugging the victim's drink.

2 shows scripts for (a) two people having drinks at a bar, and (b) kidnapping by drugging the victim's drink.

Not all situations are represented by scripts; when scripts are not available to shed light on a situation, the Response Model uses a combination of simulation and rules. Simulation is the successive application of operators that project the current state into the future until resource bounds are hit or an interesting future state is detected. Rules are one-shot predictions that result from generalization of experience that allows the recipient to jump to conclusions.

If any of these mechanisms identifies a state that is an explicit violation of an adopted character goal, the "value" – or "utility" – of that state is computed and the Response Model attempts to assess the probability of the undesirable state.

2.2 Affective Response

Affective response is computed as a function of expected utility. At this stage, the Response Model has identified a possible future state in the narrative that is deemed undesirable, e.g., has negative utility. Following [1, 2], the Response Model attempts to *solve the problem* of determining how the negative future state can be averted. That is, the Response Model computes the probability of the negative future state coming to pass as $1 - P_{\text{avert}}$, where P_{avert} is the subjective probability that the outcome will be averted either through escape bye one's own means, rescue, or happenstance (e.g., *deus ex machina*).

P_{avert} is computed using the following process: The possible future state is accompanied by a partially-ordered plan with causal links $a_i \rightarrow^p a_j$ such that action a_j could not have occurred if action a_i had not made proposition p true in the world. The plan structure is derived from the finite state machine or simulation. The Response Model attempts to intervene in the causal flow of events [6]. The Response Model uses a form of *adversarial planning* to attempt to find a new plan that negates the proposition of one of the causal links. For every action or symbol used in the new plan attempt, we compute the cost

of using that element in the plan as the difficulty of retrieval from a semantic memory model [7]. This guides the planner and also determines the total cost of each action in the new plan. The probability of averting the undesired future state is:

$$P_{\text{avert}} = (\gamma_{\text{prior}})^{tries}(\gamma_{\text{plan}})^{\sum_{i=1}^{n} \text{cost}(a_i)} \qquad (1)$$

where $tries$ is the number of causal links the Response Model tried (in order of urgency) before finding a plan that averts the possible future state, γ_{prior} and γ_{plan} are discount factors, and a_i for $i = 1...n$ are actions in the averting plan. The relative level of tension is computed as $tension = (utility)(1 - P_{\text{avert}})$ where $utility$ is the value of the possible future state that violates an adopted goal.

3 Limitations, Future Work, and Conclusions

We have limited the scope of our work to a cognitively inspired perspective of emotional response to narrative [1, 2]. Specifically, we approach the problem of computing tension as expected utility. Experimentation with the Response Model suggests that with sufficient domain knowledge (scripts, simulation operators, and rules), plausible *relative* tension results are generated. Future work includes evaluation, tuning of parameters, and expansion of system's inferencing abilities. We believe there is great value in modeling human cognitive and emotive responses to narrative. We anticipate this work will facilitate the development of story analysis systems, intelligent story authoring tools for novices, and better heuristics for story generation and interactive narrative systems.

References

1. Gerrig, R.: Experiencing Narrative Worlds: On the Psychological Activities of Reading. Yale University Press, New Haven (1993)
2. Gerrig, R., Bernardo, D.: Readers as problem-solvers in the experience of suspense. Poetics 22, 459–472 (1994)
3. Ohler, P., Nieding, G.: Cognitive modeling of suspense-inducing structures in narrative films. In: Vorderer, P., Wulff, H., Friedrichsen, M. (eds.) Suspense: conceptualizations, theoretical analyses, and empirical explorations. pp. 129–145. Lawrence Erlbaum Associates (1996)
4. Vorderer, P.: Toward a psychological theory of suspense. In: Vorderer, P., Wulff, H., Friedrichsen, M. (eds.) Suspense: Conceptualizations, Theoretical Analysis, and Empirical Explorations. pp. 233–254. Lawrence Erlbaum Associates (1996)
5. Schank, R., Abelson, R.: Scripts Plans Goals and Understanding: An Inquiry into Human Knowledge Structures. Lawrence Erlbaum Associates (1977)
6. Edwards, B., Burnett, R., Keil, F.: Structural determinants of interventions on causal systems. In: Proceedings of the 30th Annual Meeting of the Cognitive Science Society, pp. 1138–1143. Cognitive Science Society, Inc. (2008)
7. MacLeod, C., Campbell, L.: Memory accessibility and probability judgments: An experimental evaluation of the availability heuristic. Journal of Personality and Social Psychology 63, 890–902 (1992)