

Spring 5-19-2017

Predicting User Choices in Interactive Narratives using Indexter's Pairwise Event Saliency Hypothesis

Rachelyn Farrell
University of New Orleans, rfarrell@uno.edu

Follow this and additional works at: <https://scholarworks.uno.edu/td>



Part of the [Artificial Intelligence and Robotics Commons](#)

Recommended Citation

Farrell, Rachelyn, "Predicting User Choices in Interactive Narratives using Indexter's Pairwise Event Saliency Hypothesis" (2017). *University of New Orleans Theses and Dissertations*. 2319.
<https://scholarworks.uno.edu/td/2319>

This Thesis is protected by copyright and/or related rights. It has been brought to you by ScholarWorks@UNO with permission from the rights-holder(s). You are free to use this Thesis in any way that is permitted by the copyright and related rights legislation that applies to your use. For other uses you need to obtain permission from the rights-holder(s) directly, unless additional rights are indicated by a Creative Commons license in the record and/or on the work itself.

This Thesis has been accepted for inclusion in University of New Orleans Theses and Dissertations by an authorized administrator of ScholarWorks@UNO. For more information, please contact scholarworks@uno.edu.

Predicting User Choices in Interactive Narratives using Indexer's Pairwise Event Salience Hypothesis

A Thesis

Submitted to the Graduate Faculty of the
University of New Orleans
in partial fulfillment of the
requirements for the degree of

Master of Science
in
Computer Science

by

Rachelyn Farrell

B.S. University of Mississippi, 2012

May, 2017

Table of Contents

List of Figures	iii
Abstract	iv
Chapter 1: Introduction	1
Chapter 2: Related Work	3
2.1: The Indexter Model	4
2.2: Pairwise Event Salience	7
Chapter 3: Methodology	9
Chapter 4: Results	13
Chapter 5: Discussion and Future Work	15
5.1a: Referents determined by Author (Linear Story).....	16
5.1b: Referents determined by Author (Interactive Story)	17
5.2: Referents determined by Low Agency Choices	18
5.3: Proposed: Referents determined by High Agency Choices.....	19
Bibliography	21
Vita	23

List of Tables

Table 1.....	13
Table 2.....	16
Table 3.....	17
Table 4.....	18

Abstract

Indexer is a plan-based model of narrative that incorporates cognitive scientific theories about the salience—or prominence in memory—of narrative events. A pair of *Indexer* events can share up to five indices with one another: *protagonist*, *time*, *space*, *causality*, and *intentionality*. The pairwise event salience hypothesis states that when a past event shares one or more of these indices with the most recently narrated event, that past event is more salient, or easier to recall, than an event which shares none of them. In this study we demonstrate that we can predict user choices based on the salience of past events. Specifically, we investigate the hypothesis that when users are given a choice between two events in an interactive narrative, they are more likely to choose the one which makes the previous events in the story more salient according to this theory.

Keywords: *Indexer*, computational models of narrative, salience, planning, artificial intelligence, interactive storytelling

Chapter 1 - Introduction

The art of storytelling involves careful consideration for how the audience experiences the story. Skilled narrative authors pay close attention to what the audience is likely to remember from previous events and expect from future events. They may intentionally narrate events that either remind the audience of a prior event or distract them from it. This enables better control over the audience's experience and can facilitate certain discourse phenomena such as suspense and surprise. The ease with which the audience can recall a given past event—known as the *salience* of that event—is therefore a useful construct for computational models of narrative. By reasoning about the salience of events as they are narrated, narrative generation systems can leverage the same insight to achieve their intended discourse effects.

Prior research into event salience resulted in the *Indexter* model (Cardona-Rivera, Cassell, Ware, & Young, 2012), which incorporates a set of features identified by cognitive science research into a plan-based computational model of narrative that measures the salience of events according to those features. In this model, events can share up to five “narrative situation indices” with one another: *protagonist*, *time*, *space*, *causality*, and *intentionality*. That is, an event in a story can be classified according to **who** takes the action, **when** it takes place, **where** it takes place, **how** it became possible (and/or what becomes possible as a result of it), and **why** the acting characters are motivated to take the action. Using the *Indexter* model, we can begin to describe the salience of an event in terms of the situation indices it shares with other events.

A previous study (Kives, Ware, & Baker, 2015) confirmed the *pairwise event salience hypothesis*, which simply states that a past event is more salient if it shares one or more of these indices with the most recently narrated event. For example, if a past event took place in the same room as the most recently narrated event (sharing the *space* index), then that past event is easier to recall than it would

have been had it taken place in a different room. We now utilize this theory about the salience of past events to reason about the audience's expectation of the *future*. We apply this in the context of interactive storytelling, in which predicting and influencing user choices is a common research problem. Specifically, we investigate the hypothesis that when readers of an interactive narrative are given a choice between two future events, they are more likely to choose the one which will share a greater number of *Indexter* indices in total with previous events in the story—that is, the future event which will make the past more salient.

Participants read an interactive narrative, during which they made several binary choices, and were then prompted to choose between two possible endings. The number of indices that each of the two endings shared with previous events in the story depended on the prior choices the user had made. We hypothesized that they would choose whichever ending shared more indices with past events in their version of the story.

Chapter 2 - Related Work

Interactive narrative systems face an inherent tradeoff between *player agency* (the ability of the player to make meaningful choices) and *author control* (the ability of the author to control the quality of the narrative). Related research has focused on *influencing* users to make choices that are in line with the author's goal, so that the author can ensure the quality of certain branches of the story and steer users toward those branches without ultimately sacrificing player autonomy. Research toward influencing users in interactive narratives has utilized concepts from social psychology, discourse analysis, and natural language generation (Roberts & Isbell, 2014), and others have proposed lighting techniques that can be used in game environments to draw the player's attention to important elements in order to influence them to take specific actions (El-Nasr, Vasilakos, Rao, & Zupko, 2009). This study focuses on *predicting* the audience's choices rather than *influencing* them, but the implications of our findings are relevant to that area of research as well.

Indexter (Cardona-Rivera, Cassell, Ware, & Young, 2012) combines a cognitive scientific model of narrative comprehension, called the Event-Indexing Situation Model (EISM) (Zwaan & Radvansky, 1998), with a plan-based computational model of narrative. EISM is the result of decades of empirical research on how audiences store and retrieve narrative information in short-term memory while experiencing a story. Zwaan and Radvansky (1998) identify five important dimensions, or indices, of narrative events which have been shown to play a role in narrative comprehension: *protagonist* (who), *time* (when), *space* (where), *causality* (how), and *intentionality* (why). *Indexter* defines a plan data structure augmented with this model. The story is divided into a series of discrete events, and at each moment *Indexter* measures the salience of each past event.

Numerous plan-based models have been used to reason about story structure and to control interactive stories (see survey by Young et al. (2013)). Plan-based models have also been used to achieve

other discourse phenomena, such as suspense (Cheong & Young, 2008), surprise (Bae & Young, 2014), and cinematic composition (Jhala & Young, 2010). As with these other models of discourse, Indexter can inform story generation as well as discourse generation.

Indexter has also been used to predict agency in interactive stories (Cardona-Rivera, et al., 2014). When choosing between two alternatives in a hypertext adventure game, players self-reported a higher sense of agency when the perceived next state that would result from making each choice differed from one another in at least one index. This study and (Kives, Ware, & Baker, 2015) suggest that *Indexter* might be used not only to measure the salience of past events but also the degree to which the audience expects future events—what Young and Cardona-Rivera (2011) call a *narrative affordance*. Recent work by these researchers (Cardona-Rivera & Young, 2014) has explored a more nuanced model of narrative memory, but we demonstrate that interesting results can be obtained even with the simple pairwise event salience model.

2.1 - The Indexter Model

Indexter defines a data structure for representing stories as plans. Under the pairwise event salience model, a pair of events in a story can share up to five dimensions with one another:

protagonist, time, space, causality, and intentionality. This section reproduces very briefly those definitions needed to understand the evaluation described in this paper; for a detailed description of how Indexter maps EISM indices to plan structures, see the description by Cardona-Rivera et al. (2012)

A plan is a sequence of events that achieves some goal (Russell & Norvig, 2010). Each event has preconditions which must be true immediately before it is executed, and effects which modify the world state immediately afterwards. The kinds of events that can occur in a given domain are represented by abstract, parameterized templates called operators, as described by the STRIPS formalism (Fikes &

Nilsson, 1972). Beginning with the initial world state, a planner finds a sequence of events that, when executed in order, modify the world such that the goal is achieved at the end.

Narrative planning (Riedl & Young, 2010) is an application of planning in which the domain is a story world, the goal is the author's goal for the end state of the story, and the operators represent possible story events. For example, the domain might define an operator *attack(?attacker*, ?victim, ?location)*. Each term starting with a '?' is a free variable which can be bound to a constant corresponding to some specific thing defined in the story world. The preconditions might be that the attacker and victim are both alive, that both are in the same location, that the attacker is armed, and that the victim is unarmed. The effects that take place as a result of the *attack* event might be that the victim is no longer alive.

In addition to the author's goal, a narrative plan reasons about a different type of goal: the goals of the individual characters. The planner does not attempt to achieve the characters' goals; instead they are used simply to *explain* the actions taken by the characters. A narrative plan must satisfy the constraint that all actions included in the plan are explained for all characters who take them. To achieve this, each event template must specify whether any of its parameters represent *consenting characters*, or characters who are responsible for taking that action. In the *attack* example, the *?attacker* is the sole consenting character (marked with an asterisk), because this is the only character who must be willing to carry out the attack. (While the *?victim* is also a character involved in the action, it need not *consent* to being attacked.) An action is said to be *explained* for a character if there is some possible branch of the story in which that action contributes to one or more of that character's goals.

Each operator in an *Indexer* plan must also specify two required parameters: *?time* (the time frame in which it occurs) and *?location* (the location at which it occurs). We can now define an *Indexer* event as a fully ground instance of such an operator. For example, an attack event might be:

attack(?attacker=Roy, ?victim=Dirk, ?location=Gym, ?time=Day2). The following definitions are used to determine whether two *Indexter* events share each of the EISM indices.

Definition 1. Two events share the *protagonist* index iff they have one or more consenting characters in common.¹

Definition 2. Two events share the *time* index iff their time parameters are the same symbol.

Definition 3. Two events share the *space* index iff their location parameters are the same symbol.

Cognitive science research (Magliano, Miller, & Zwaan, 2001; Zacks, Speer, & Reynolds, 2009) has demonstrated that time and space can be hierarchically organized in memory. Whether different rooms in the same house count as different locations depends on the discourse. *Indexter* uses a simplified representation of these concepts as unique symbols. For this to be effective, the appropriate level of granularity must be communicated to the audience.

One strength of the plan-based models of narrative on which *Indexter* is based is the ability to reason about causal relationships between events. While cognitive scientists have studied several forms of causality (Trabasso & Sperry, 1985; Trabasso & Van Den Broek, 1985; Zwaan & Radvansky, 1998), one in particular is easily available in plans using causal links: the ways in which the effects of earlier events enable later events by establishing their preconditions.

Definition 4. A causal link $s \xrightarrow{p} t$ exists from event *s* to event *t* for proposition *p* iff *s* occurs before *t*, *s* has the effect *p*, *t* has a precondition *p*, and no event occurs between *s* and *t* which has the effect $\neg p$. We say that *s* is the causal parent of *t*, and that an event's causal ancestors are those events in the transitive closure of this relationship.

¹ Here we use the one protagonist per event (as opposed to one per story) definition discussed by Cardona-Rivera et al. (2012).

Definition 5. Two events share the *causality* index iff the earlier event is the causal ancestor of the later event.

Riedl and Young's intentional planning framework organizes events into *frames of commitment* to explain how characters achieve their individual goals. These structures also rely on consenting characters and causal relationships.

Definition 6. Two events share the *intentionality* index iff both events can be explained for some character by the same goal; that is, both events have *c* as a consenting character, *c* has a goal *g*, and both events are causal ancestors of some event that has the effect *g*. In other words, two events share intentionality when both are taken by the same character for the same goal.

2.2 – Pairwise Event Salience

The *pairwise event salience hypothesis* was proposed in the original description of *Indexter* as a starting point for a model of how narrative situation indices are correlated to salience. The authors suggested a series of studies that begin by evaluating this model in its simplest form, and then iteratively incorporate new insights to build up a more powerful model. A recent study (Kives and Ware, 2015) began this process, defining the model as follows: When a past event shares one or more indices with the most recently narrated event, that past event is more salient than one which shares no indices with the most recently narrated event.

These authors conducted a study which evaluated the pairwise event salience hypothesis using reaction time as a proxy for salience. In this experiment, participants read short text stories one event at a time. A certain event in each story was designated the *referent*, and a later event was designated the *prompt*. The stories were designed such that the prompt and referent shared exactly one or zero

indices. After reading the prompt, participants were interrupted and asked to recall the referent, and the speed with which they were able to answer was used as a proxy to measure salience.

This experiment tested five hypotheses individually—one for each of the five indices. The hypothesis for a given index was that subjects would react faster when the prompt shared that index with the referent than when it did not share any indices. Paired t-tests were used to compare each index to the None condition. Of the five null hypotheses, four were rejected at the $p < 0.05$ level, while one (*causality*) could only be rejected at the $p < 0.1$ level. The authors concluded that participants who accurately remembered the referent were able to respond faster when the referent shared at least one index with the most recently narrated event.

The results of this study provide support for the use of *Indexter* to measure the salience of past events. The present study builds upon this notion by demonstrating that the audience's desires and expectations for *future* events are affected by the salience of past events, and thus *Indexter* can be an effective tool in reasoning about the audience's mental model of the future.

Chapter 3 - Methodology

We designed an interactive story whose events could be generated by a narrative planner. Readers must choose between two possible endings to the story, and we hypothesized that they would choose the ending whose *Indexter* event shared more indices with previous events in the story. We allowed the user to make four intermediate choices throughout the story, each of which determined the value for a specific index of a single event. The four choices tested the *protagonist*, *time*, *space*, and *intentionality* indices, respectively. (We chose to exclude the *causality* index due to the added complexity of including a choice which toggles between two events, each of which is causally related to only one ending, while still ensuring that both endings are possible in all versions of the story.) There were two possible values for each index tested, thus a total of 16 possible story configurations. When the reader arrived at the final choice, the number of indices that each of the two possible ending events shared with the rest of the story was determined solely by their four previous choices.

In the story, two prisoners are threatened to be killed by the prison bully, and each of them comes up with a different plan in response. Ernest plans to break out of prison and escape onto the highway, and Roy plans to kill the bully in the gym. Both plans involve stealing an item and then crawling into the ductwork through a loose vent. In all versions of the story, both characters end up inside the ductwork, ready to complete one of the two plans together; but a guard discovers their whereabouts at the last minute. However, the guard believes there is only one prisoner in the duct, not two. Roy and Ernest realize that if they continue, they will both be caught and neither goal will be accomplished; but if one of them turns himself in, the other can still proceed with his original plan. The user's final choice is between the "Escape" ending (where Ernest escapes onto the highway) and the "Revenge" ending (where Roy kills the bully in the gym). The following is a description of how we manipulate each *Indexter* index before arriving at this experimental choice.

Protagonist: The story begins with the two prisoners discovering a hidden pack of cigarettes which turns out to belong to the prison bully. This angers the bully, who threatens to kill them both. The user makes the seemingly arbitrary choice of which character takes the cigarettes. The chosen character will later be given an extra scene by himself; after being caught by a guard while stealing his item, that character must complete a punishment duty. The additional scene of this character fulfilling his punishment introduces a new event into the story which shares the *protagonist* index with that character's ending but not the other.

Space: For the punishment scene, the user chooses between two possible punishments for that character to do—picking up trash off the highway, or cleaning the equipment in the gym. This introduces an additional event matching the *space* index of one of the two endings, since the escape ending will take place on the highway, and the revenge ending will take place in the gym. To communicate the appropriate level of granularity for the space index, we displayed a graphic on each passage showing the layout of the prison with the location of the current event highlighted and labeled, e.g. “highway”, “gym”, “cafeteria”, etc.

Time: Since both endings have the same symbol for *time* (Day 2), we cannot simply add a new event that shares the *time* index with one ending but not the other. Therefore to account for the *time* index we must deviate slightly from our pattern. The two theft scenes—Ernest stealing some disguises for his escape plan, and Roy stealing a knife for his revenge plan—are told in variable order depending on the user's choice; one takes place on Day 1, and the other on Day 2. We expect that whichever theft scene happens on Day 2 will be more salient when the final choice is made, and therefore the user will be more likely to choose the ending involving the same protagonist and intentionality as that theft scene. In other words, if the most recent theft event was Roy stealing the knife for the goal of revenge, then the *time* vote goes to Roy's revenge ending. To communicate the granularity of the time index, we

displayed a graphic of a calendar on each passage, showing either “Day 1” or “Day 2” according to the time of the current event.

Intentionality: After the second theft is completed on Day 2, the user chooses between two preparatory actions for both characters to take together: either donning the disguises for the goal of escape—which introduces a new event sharing the *intentionality* index with the escape ending—or locking the bully in the gym, which does the same for the revenge ending.

Next, the characters take the necessary step of sneaking into the air duct—from which they plan to exit either into the gym where they can kill the bully, or outside where they can escape on the highway. Finally, the guard catches up to them and we prompt the user for the final choice.

To summarize: If the escape ending is chosen, the final event will have the parameters (character=Ernest, location=highway, time=Day 2, goal=escape). If the revenge ending is chosen, it will have the parameters (character=Roy, location=gym, time=Day 2, goal=revenge). The hypothesis is that the user will choose the ending for which more of the following are true:

- Its character is the same as the character who had one extra scene (protagonist)
- Its location is the same as the location of the punishment scene (space)
- Its character is the same as the character who stole his item on Day 2 (time)
- Its goal is the same as the goal of the preparatory action (intentionality)

We built the story using Twine, an open-source tool for writing branching stories. We recruited 350 participants through Amazon Mechanical Turk, and paid them each \$0.25 for completing the story. To adjust for the high volume of noise on Mechanical Turk, we asked each user a series of comprehension questions after they completed the story. The questions were designed to verify that the story accurately communicated the pertinent information to the user. Each question displayed two events from the version of the story they read—one from the ending scene and one from a previous scene—and asked a question such as, “Were these two actions taken by the same character?” or “Did these two

events happen in the same place?” We discarded the data from participants who did not answer all of the comprehension questions correctly, and gave an additional \$0.75 bonus to those who did.

Participants were aware of the available bonus from the start.

Chapter 4 - Results

Of the 350 results, we discarded 225 and were left with 125 responses from participants who demonstrated full comprehension of the story. Because we were not attempting to influence readers to choose one path or the other, users were free to make exactly two choices in favor of the Escape ending and two in favor of the Revenge ending. In these cases, we make no prediction as to which ending they would choose. Of the remaining 125 results, there were 78 for which a majority of the user's choices were in favor of one ending or the other. We conducted the following evaluation using those 78 results.

To evaluate our hypothesis we used Fisher's exact test, which is similar to the χ^2 test but performs better for distributions with small expected values (Fleiss, Levin, & Paik, 2013). Fisher's exact test is nonparametric, meaning it does not assume any underlying distribution of the population. This is important because participants chose more Escape options overall than Revenge ones (perhaps due to the morality differences between the two paths). Fisher's exact test is not skewed by this imbalance. Table 1 shows the contingency table giving the frequency distribution of results according to their expected outcomes.

Table 1. Contingency table for Fisher's Exact Test

	Chose <i>Escape</i>	Chose <i>Revenge</i>
Expected <i>Escape</i>	32	14
Expected <i>Revenge</i>	11	21

The null hypothesis was that the ending choices were independent of the *Indexter* indices of previous events. Fisher's exact test rejected this with $p < 0.0022$. There are several ways to measure effect size when using Fisher's exact test. The odds ratio for this contingency table is ≈ 4.27 , meaning

there are about 4 to 1 odds that users chose the ending we expected them to choose. We conclude that users are indeed more likely to choose future events which will make past events more salient.

Chapter 5 - Discussion and Future Work

We have demonstrated that interactive narrative systems can make use of *Indexter* indices to predict user choices, even using the simple *pairwise event salience* model. We have shown that, when presented with choices for future events, users significantly prefer those events which have more indices in common with prior events in the story. We believe that plan-based narrative systems can utilize this information about the audience's desires and expectations both to reason about discourse phenomena such as suspense and surprise, as well as to influence users to make choices that are in line with the author's intentions.

Toward the goal of *influence*, we have some additional observations to make. These results may seem to suggest that we could influence users simply by inserting events into the story that share indices with the later events we want them to choose. However, in doing so we would be altering one potentially crucial element of the structure used in this study: the element of *choice*.

In this study, the past events that acted as *referents* (the events that could share indices with one of the two endings) were directly chosen by the user. Furthermore, those choices were *high agency* choices according to a definition made by Cardona-Rivera et al. (2014), meaning that the options differed by at least one index. It is unclear from our study alone whether users' ending choices could be influenced by the indices of *any* prior event, or specifically by the indices of prior events that they *chose* to happen, or even more specifically, by the indices of prior events that resulted from *high agency* choices.

Since the completion of this study, we have conducted a series of follow-up experiments to examine these alternatives more closely. The following section describes three experiments, all of which were based on the same story used here, but with modifications designed to target the different

scenarios outlined above. Lastly, I discuss the insight gained from this series and propose a fourth and final installment.

5.1a – Referents determined by Author (Linear Story)

The simplest modification of this study that would enable us to influence users’ ending choices is to predetermine the events that make one ending more salient—that is, instead of allowing the user to make *choices* that determine the referent events, we simply choose for them. If the element of choice is inconsequential to the results found in the original study, then we should expect to be able to influence users to choose the *Escape* ending, for example, simply by selecting all of the choices that make the *Escape* ending more salient.

We created two versions of the story in this manner; one in which we selected the Escape-salient option for all four preliminary choices, and one in which we selected all four Revenge-salient options. We removed the prompts for those choices and simply narrated the resulting linear story. In these versions, the users made no choices other than the experimental choice at the end.

We divided participants into two groups; one for each story version. Using the same experimental setup as the original study, we had 32 users read the Escape-salient version and correctly answer validation questions, and 36 users for the Revenge-salient version. The results are presented in Table 2.

Table 2. Contingency table, Referents determined by Author (Linear Story)

	Chose <i>Escape</i>	Chose <i>Revenge</i>
Expected <i>Escape</i>	17	15
Expected <i>Revenge</i>	22	14

Fisher’s exact test fails to reject the null hypothesis, giving a p-value of 0.8286. We conclude that the element of choice is *not* entirely inconsequential.

5.1b – Referents determined by Author (Interactive Story)

In the previous version, choices were removed entirely. There is a possibility that the failure of that version was due to the lack of interactivity of the story in general. Thus, for the next experiment we made a simple adaptation; we added choices that affected only events that were *unrelated* to the endings. None of these choices shared any indices with either of the two endings; the four referents were still predetermined by the author. In the same manner as the previous version, we created two versions of the story and attempted to influence users in each group to choose the associated ending. We had 48 results from the Escape version, and 40 for the Revenge version. The results are shown in Table 3.

Table 3. Contingency table, Referents determined by Author (Interactive Story)

	Chose <i>Escape</i>	Chose <i>Revenge</i>
Expected <i>Escape</i>	27	21
Expected <i>Revenge</i>	26	14

Fisher’s Exact Test fails once again to reject our null hypothesis, with a p-value of 0.854. We conclude that the failure of the previous version was *not* simply due to the lack of choices present in the story. We believe that it was instead due to the lack of choices relating to the endings—that is, the *referent* events were not the results of choices. We believe that in order to successfully influence users by manipulating *Indexter* indices of events, we must consider specifically the indices of events that the user *chose* to happen. The next approach builds on this insight.

5.2 – Referents determined by Low Agency Choices

In this version, we allowed users to make choices for the referent events, as was the case in the original study. However, unlike the original, in this version both choices shared the *same value* for the target index. For example, users in the Escape group were prompted to choose between two punishments: 1) Picking up trash along the highway, or 2) Mowing the grass along the highway. In either case, the location of the resulting event is the Highway, and thus it will share the *space* index with the Escape ending. Similarly, users in the Revenge group were given a choice between two punishment duties that both take place in the Gym.

We again created two versions of the story where all four choices were handled in this way; either all sharing indices with the Escape ending, or all sharing indices with the Revenge ending. We had 36 valid responses from the Escape group and 38 from the Revenge group. As shown in Table 4, we were once again unsuccessful in influencing users' final choices. Fisher's Exact Test fails to reject our null hypothesis, with a p-value of 0.7451.

Table 4. Contingency table, Referents determined by Low Agency Choices

	Chose <i>Escape</i>	Chose <i>Revenge</i>
Expected <i>Escape</i>	22	14
Expected <i>Revenge</i>	25	13

From this, we make an important observation about the role of *agency* in predicting and influencing choices. These follow-up studies suggest that our success in predicting users' ending choices in the original study was contingent on the fact that the referent events—the events that were made salient by the ending events—were direct results of *high agency* choices. To summarize, we believe that

when readers of interactive stories are given *meaningful choices*, they have a tendency to choose future events which make their past choices more salient.

5.3 – Proposed: Referents determined by High Agency Choices

In the spirit of completion, I have proposed one final addition to this series, in which we will once again attempt to influence users' choices based on the *pairwise event salience* model, this time ensuring that the referent events are direct results of *high agency* choices. A simple modification of the previous version will suffice: For a given choice, rather than making both options share the *same* value for the target index, make only *one* option have that value, and the other option have some *other* value that is not shared by either ending.

For example, the punishment options for the Escape group might be: 1) Picking up trash along the highway, or 2) Cleaning the facilities in the bathroom. This is a high agency choice because the events differ in at least one index (*space*), but it allows only one ending (Escape) to potentially make this event more salient in the end. If all four choices are designed in this manner, then it is most likely that users in this group will choose *at least one* event that shares an index with the Escape ending. Specifically, there is only a 1/16 chance that they will choose all four of the *other* options; in the remaining 15/16 cases, by the time the prompt is reached, the Escape ending will make the users' prior choices more salient than the Revenge ending. Based on our conclusions from the previous experiments, this version should finally succeed in influencing users to choose the target ending.

If this attempt fails, it could mean that the only way to predict choices using the simple *pairwise event salience* model is by giving users actual freedom to determine which ending makes their choices more salient. In other words, the referent events would not only need to be results of high agency choices, but specifically choices that could be co-salient with *either* of the two endings, as was the case in our original study. If this is true, then *influencing* the audience using this model alone would be

substantially more difficult, and therefore less useful in practice. In this case, we would conclude that the *pairwise event salience* model is insufficient to enable influence. Indeed, regardless of the outcome of the proposed experiment, we will use the insight gained from all of these studies to develop a more detailed model of salience, as discussed in previous sections, and improve our understanding of how agency affects this model in the interactive context.

Bibliography

- Bae, B.-C., & Young, R. (2014). A computational model of narrative generation for surprise arousal. *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games*, 6, 131-143.
- Cardona-Rivera, R. E., & Young, R. M. (2014). A Knowledge Representation that Models Memory in Narrative Comprehension. *Proceedings of the 28th AAAI Conference on Artificial Intelligence - Student Abstracts Track*, (pp. 3098--3099).
- Cardona-Rivera, R. E., Cassell, B. A., Ware, S. G., & Young, R. M. (2012). Indexter: a computational model of the Event-Indexing Situation Model for characterizing narratives. *Proceedings of the 3rd Workshop on Computational Models of Narrative*, (pp. 34--43).
- Cardona-Rivera, R. E., Robertson, J., Ware, S. G., Harrison, B., Roberts, D. L., & Young, R. M. (2014). Foreseeing meaningful choices. *Proceedings of the 10th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, (pp. 9--15).
- Cheong, Y.-G., & Young, R. M. (2008). Narrative generation for suspense: modeling and evaluation. *1st International Conference on Interactive Digital Storytelling*, (pp. 144--155).
- El-Nasr, M. S., Vasilakos, A. V., Rao, C., & Zupko, J. A. (2009). Dynamic Intelligent Lighting for Directing Visual Attention in Interactive 3-D Scenes. *IEEE Trans. Comput. Intellig. and AI in Games*, 145--153.
- Fikes, R. E., & Nilsson, N. J. (1972). STRIPS: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2, 189--208.
- Fleiss, J., Levin, B., & Paik, M. (2013). *Statistical methods for rates and proportions*. Wiley.
- Jhala, A., & Young, R. M. (2010). Cinematic visual discourse: representation, generation, and evaluation. *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games*, 2, 69--81.
- Kives, C., Ware, S. G., & Baker, L. J. (2015). Evaluating the pairwise event salience hypothesis in Indexter. *Proceedings of the 11th AAAI International Conference on Artificial Intelligence and Interactive Digital Entertainment*, (pp. 30--36).
- Magliano, J. P., Miller, J., & Zwaan, R. A. (2001). Indexing space and time in film understanding. *Applied Cognitive Psychology*, 15, 533--545.
- Riedl, M. O., & Young, R. M. (2010). Narrative planning: balancing plot and character. *Journal of Artificial Intelligence Research*, 39, 217--268.
- Roberts, D. L., & Isbell, C. L. (2014). Lessons on Using Computationally Generated Influence for Shaping Narrative Experiences. *IEEE Transactions on Computational Intelligence and AI in Games*, 6, 188-202.
- Russell, S., & Norvig, P. (2010). *Artificial Intelligence: a modern approach* (Third ed.). Prentice Hall.
- Trabasso, T., & Sperry, L. L. (1985). Causal relatedness and importance of story events. *Journal of Memory and language*, 24, 595--611.

- Trabasso, T., & Van Den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of memory and language, 24*, 612--630.
- Young, R. M., & Cardona-Rivera, R. E. (2011). Approaching a Player Model of Game Story Comprehension Through Affordance in Interactive Narrative. *Proceedings of the 4th Workshop on Intelligent Narrative Technologies*, (pp. 123--130).
- Young, R. M., Ware, S. G., Cassell, B. A., & Robertson, J. (2013). Plans and planning in narrative generation: a review of plan-based approaches to the generation of story, discourse and interactivity in narratives. *Sprache und Datenverarbeitung, Special Issue on Formal and Computational Models of Narrative, 37*, 41--64.
- Zacks, J. M., Speer, N. K., & Reynolds, J. R. (2009). Segmentation in reading and film comprehension. *Journal of Experimental Psychology: General, 138*, 307.
- Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin, 123*, 162.

VITA

The author was born in Durham, North Carolina and grew up in Oxford, Mississippi. She obtained her Bachelor's degree in computer science from the University of Mississippi in 2012. She joined the University of New Orleans computer science graduate program in 2014, and became a member of the Narrative Intelligence Lab in 2015.