Proceedings, The Thirteenth AAAI Conference on
Artificial Intelligence and Interactive Digital Entertainment (AIIDE-17)

Influencing User Choices in Interactive Narratives Using Indexter's Pairwise Event Salience Hypothesis

Rachelyn Farrell, Stephen G. Ware
rfarrell@uno.edu, sgware@uno.edu
Narrative Intelligence Lab, Department of Computer Science
University of New Orleans, New Orleans, LA, 70148, USA

Abstract

Indexter 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 Indexter events can share up to five indices with one another: *protagonist, time, space, causality,* and *intentionality.* The pairwise event salience hypothesis states that a past event is more salient if it shares one or more of these indices with the most recently narrated event. In a previous study we used this model to predict users' choices in an interactive story based on the indices of prior events. We now show that we can use the same method to *influence* them to make certain choices. In this study, participants read an interactive story with two possible endings. We influenced them to choose a particular ending by manipulating the salience of story events. We showed that users significantly favored the targeted ending.

Introduction

Human storytellers often pay close attention to how they can affect the audience's memory of certain story events. They can intentionally narrate events that make the audience more likely to remember or forget things that previously happened. This can be used to achieve certain narrative discourse effects like surprise and suspense. The ease with which a person can recall a given past event is known as the *salience* of that event.

The Event-Indexing Situation Model (EISM) is a cognitive scientific model that identifies five indices by which narrative events are stored and retrieved in short-term memory—protagonist (who), time (when), space (where), causality (how or what enabled), and intentionality (why) (Zwaan and Radvansky 1998). Indexter incorporates these indices into a plan-based model of narrative in order to reason about the salience of narrative events (Cardona-Rivera et al. 2012).

The original description of Indexter proposed a starting point for how to calculate the salience of a narrative event based on the indices of the current (or most recently narrated) event. This was later formalized and experimentally validated as the *pairwise event salience hypothesis*, which states that a past event is more salient if it shares at least one index with the most recently narrated event (Kives, Ware, and Baker 2015). We subsequently showed that the salience of past events can be used to predict readers' choices for *future* events, even based on the simple pairwise model (Farrell and Ware 2016). Participants in that study were given two possible endings, and were significantly more likely to choose the ending which made the events that resulted from their *past choices* more salient.

Now, we demonstrate that the same method can be applied to *influence* readers to choose a particular ending. We developed an interactive story with two possible endings, and engineered the available choices throughout the story such that one ending would make the reader's past choices more salient, and the other would not. Users were significantly more likely to choose the ending that we targeted. Furthermore, they reported that they did not feel influenced by the author to choose either ending.

Related Work

Interactive narrative systems face a 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). Researchers have studied various approaches to 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. Roberts and Isbell (2014) utilized concepts from social psychology, discourse analysis, and natural language generation to influence users. El-Nasr et al. (2009) 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. Our study approaches the challenge of influence by reasoning about the audience's memory of previous events.

Indexter (Cardona-Rivera et al. 2012) is a computational cognitive model that reasons about how the audience comprehends a story's *discourse*, or how the story is told (Bal 1997). Indexter defines a plan data structure based on IPOCL (intentional partial-order causal link) plans (Riedl and Young 2010), augmented with a cognitive scientific model of narrative comprehension called the event-indexing situation model, or EISM (Zwaan and Radvansky 1998). EISM is the result of decades of empirical research on how audiences store and retrieve narrative information in short-term memory while experiencing a narrative. 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 or what enabled), and intentionality (why).

Plan-based models have been applied to other discourse phenomena, such as suspense (Cheong and Young 2008), surprise (Bae and Young 2014), and cinematic composition (Jhala and Young 2010). Numerous plan-based models have been used to reason about story structure and to control interactive stories (see survey by Young et al. 2013). 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 suggests that Indexter might be used to measure not only the salience of past events, but also the degree to which the audience expects future events—called a *narrative affordance* (Young and Cardona-Rivera 2011). Recent work by these researchers (Cardona-Rivera and 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.

The Indexter Model

Indexter defines a data structure for representing stories as plans. 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. (Cardona-Rivera et al. 2012).

The story is divided into a series of discrete events, and at each moment Indexter measures the salience of each past event as a function of the indices it shares with the current event. A pair of events in a story can share up to five dimensions with one another: protagonist, time, space, causality, and intentionality. Events are represented using abstract, parameterized templates, or operators, as described by the STRIPS formalism (Fikes and Nilsson 1972). For example, the domain might define an operator *attack(?attacker*, ?victim, ?time, ?place)*. Each term starting with a '?' is a free variable which can be bound to a constant corresponding to some specific thing in the story world.

Each event in an Indexter plan must specify two required parameters: the time frame in which it occurs and the location at which it occurs. For example, the *attack* action might be instantiated with *?time* = Day2 and *?place* = Gym. Cognitive science research (Magliano, Miller, and Zwaan 2001; Zacks, Speer, and 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 and Sperry 1985; Trabasso and Van Den Broek 1985; Zwaan and Radvansky 1998; Porteous et al. 2017), 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 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 s is the *causal parent* of t, and that an event's *causal ancestors* are those events in the transitive closure of this relationship.

Riedl and Young's intentional planning framework, upon which Indexter is based, organizes events into *frames of commitment* to explain how characters achieve their individual goals (Riedl and Young 2010). An operator may specify any number of its parameters as being the *consenting characters*, or those who must consent to the action before it can be included in the plan. In the previous example, the *?attacker** parameter represents the sole consenting character (denoted with an asterisk) because it is solely responsible for carrying out the attack. While the *?victim* is clearly a character involved in the action, it need not be a consenting character. A frame of commitment is the series of actions taken by some character *c* in service of some goal *g*, beginning with the motivating step (the step which caused *c* to adopt the goal *g*), and ending with the final step, whose effects finally achieve *g*. All other steps in the frame must be causal ancestors of the final step, and have *c* as a consenting character.

Using these definitions we can now formalize the Indexter model as follows.

1. Two events share the *protagonist* index iff they have one or more consenting characters in common.[1]
2. Two events share the *time* index iff their time parameters are the same symbol.
3. Two events share the *space* index iff their location parameters are the same symbol.
4. Two events share the *causality* index iff the earlier event is the causal ancestor of the later event.
5. Two events share the *intentionality* index iff they are part of the same frame of commitment—in other words, if they are part of the same character plan.

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 (Cardona-Rivera et al. 2012). A later study (Kives, Ware, and Baker 2015) defined the model as follows: When

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

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. The authors conducted a study to evaluate the hypothesis using readers' reaction time as a proxy for salience. They demonstrated that when readers were interrupted during a story and asked to recall past events, those who remembered the events accurately were able to remember them faster when they shared at least one index with the most recently narrated event. This supports the use of Indexter to measure the salience of past events.

Our previous study (Farrell and Ware 2016) built upon that notion by demonstrating that the audience's desires and expectations for *future* events are affected by the salience of past events. Specifically, we tested the claim that when users of an interactive narrative are given a choice between two future events, they will choose the one that makes the *past* more salient according to the pairwise model—in other words, the one that shares more indices in total with previous events. Participants read an interactive story with two possible endings, and were given choices throughout the story that caused certain events to share an index with one ending but not the other. In that study, we did not attempt to influence users to choose a particular ending—it was their choices throughout the story that determined which ending we predicted for them. In the present study, we show that the pairwise event salience model can be used not only to predict readers' choices, but also to influence them.

Methodology

For this study we used a modification of the story used in the previous study. The story is about two prisoners who are threatened by the prison bully. They each come up with a different plan in response. Ernest wants to *escape* the prison onto the highway; Roy wants to get *revenge* by killing the bully in the gym. Roy and Ernest work together in pursuit of both goals throughout the story, but at the last minute, they are forced to make a choice. One of them must turn himself in, but the other gets to carry out his plan. The reader makes the final choice between the Escape ending, where Ernest escapes onto the highway, or the Revenge ending, where Roy kills the bully in the gym:

Escape: *escape(Ernest*, Highway, Day2)*
Revenge: *attack(Roy*, Dirk, Gym, Day2)*

In the previous study, readers were given choices throughout the story for events that would share an index with either one of the two endings. The endings otherwise had the same number of each index in common with previous events in the story. We showed that if the majority of the reader's earlier choices matched indices of one ending, they were significantly more likely to choose that ending.

An important observation about that study is that we only considered events that were direct results of the reader's *choices.* We do not know if we can predict choices for future events based on their shared indices with *any* past events, only on shared indices with past events that were knowingly chosen by that reader. We conducted a series of experiments to test simpler approaches to salience-based influence which were unsuccessful. We concluded from these experiments that the aspect of choice was crucial in our previous study, and that influencing readers is not as simple as, for example, inserting several events that share indices with the target ending. We need to make the reader *choose* those events directly.

We randomly divided participants into two groups and attempted to influence one group to choose the Escape ending and the other to choose the Revenge ending. Readers were given four intermediate choices throughout the story, and each choice presented two options for what the next event will be. In all four cases, one of these events shared an index with the targeted ending, and the other did not share it with either ending. Whenever the reader chooses the event that *does* share the index, we consider this one "vote" for the targeted ending. In other words, this becomes one past choice that the targeted ending will make more salient. We can predict that the reader will choose the targeted ending as long as it receives at least one of these four potential votes, since there are no options that can give any votes to the other ending.

Each intermediate choice tests a different index. In the previous study we did not include the causality index due to the significant design challenges it would have introduced. Although it would have been feasible in this study, we chose to exclude it again for consistency. The following section describes how we manipulate the protagonist, time, space, and intentionality indices of the intermediate choices before arriving at the final experimental choice.

Protagonist: The story begins with Roy and Ernest learning that the prison bully intends to kill them. They learn this through a friendly guard, who gives them his key card in hopes that they can find a way to use it. Roy and Ernest both happen to share a chore duty with the bully, so their first priority is to get their chore duties changed somehow. The reader is given two options—either our target ending's protagonist gets some friends to swap chores with them, or the friendly guard changes their chores on the official schedule. The resulting event will therefore either share a protagonist with our target ending, or it will share no protagonists with either ending.

Time: As with the previous study, we must handle the time index differently because the two endings have the same value for that index (Day 2). Therefore any past choice which shares the time index with the Revenge ending will also share it with the Escape ending, so we cannot say that only one ending makes that event more salient. To determine a vote for the time index, we utilize the other indices. Roy's plan to kill the bully requires him to steal a knife, while Ernest's plan to escape the prison requires him to steal some civilian clothes. One of these theft scenes will happen on Day 1, but the other (the one with our target ending's protagonist) can either happen on Day 1 or Day 2, depending on the reader's choice. If they choose for it to happen on Day 2, then the ending whose protagonist performed that theft gets the time vote. Otherwise neither ending gets the vote.

Space: The character performing the second theft gets caught by a guard and sentenced to a punishment. The reader is given the choice between one punishment that shares the space index with the targeted ending (either picking up trash along the Highway for the Escape group, or wiping down the equipment in the Gym for the Revenge group), or cleaning the bathroom, which shares the space index with neither of the two endings.

Intentionality: Roy and Ernest have time to do one more thing together before they enter the final phase of their plans. They can either take some action that contributes to the same goal as the target ending (either changing into their disguises for the Escape group, or readying their knives for the Revenge group), or one that does not contribute to either goal—returning the key card to the guard who helped them.

We hypothesize that participants in the Escape group will choose the Escape ending, and participants in the Revenge group will choose the Revenge ending, except in cases where the participant made all four of the *other* choices; in those cases the two endings shared an equal number of indices with past choices, so we have no prediction to make. We built the story using Twine (Friedhoff 2013), an open-source tool for writing branching stories. We recruited 260 participants through Amazon Mechanical Turk, and paid them each $0.10 for completing the story.

To adjust for the high volume of noise on Mechanical Turk, we asked each participant a series of comprehension questions after they completed the story. The questions were designed to verify that the reader understood the specific information that was important for the study—for example, where the punishment scene took place. We gave a $0.90 bonus to participants who answered all of these questions correctly, and discarded the data of anyone who did not. Participants were made aware of the available bonus from the start.

Results

After discarding the data from participants who did not answer the comprehension questions correctly, we were left with 124 responses from participants who demonstrated full comprehension of the story and also made at least one choice that gave a vote to the targeted ending. To evaluate whether readers were significantly more likely to choose that ending, 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, and Paik 2013). It is also nonparametric, which means it does not assume any underlying distribution of the population. This is important because we are not assuming that the two endings are equally favorable. In fact, readers generally preferred the Escape ending with about 3 to 1 odds, possibly due to the morality differences between the two events. Fisher's exact test is not skewed by this imbalance.

Table 1 gives the frequency distribution of the results according to their predicted outcomes. The *p*-value given by Fisher's exact test is the likelihood that this distribution would occur if there were no association between the rows and the columns.

Table 1: Results (at least 1 vote)

	Escape Ending	Revenge Ending
Escape Group	46	16
Revenge Group	32	30

$p < .0076$

The null hypothesis, that readers' ending choices were independent of the Indexter indices of their previous choices, was rejected by Fisher's exact test ($p < .05$). The odds ratio for this table is $\simeq 2.67$, meaning that there are about 2.67 to 1 odds that the reader chose the ending we were influencing them to choose.

Note that the Revenge group was almost evenly divided between the two endings. Although it may appear that we were only successfully influencing the Escape group, the statistical result as well as our own experience from previous studies suggest that a significantly larger percentage of readers in the Revenge group would have chosen the Escape ending had we not influenced them to choose Revenge.

We performed two additional analyses considering separately those readers who made at least *two* choices in favor of the targeted ending, and those who made at least *three*. We expected to see an increasingly stronger preference for the targeted ending as we filter the responses in this way. Of the 124 responses used in the first analysis, there were 69 for which the reader made at least two choices in favor of the target ending. Table 2 shows the frequency distribution of this subset.

Table 2: Subset of Results (at least 2 votes)

	Escape Ending	Revenge Ending
Escape Group	30	10
Revenge Group	15	14

$p < .040$

In this case, the null hypothesis was rejected again, this time with an odds ratio of 2.76. Finally, Table 3 shows the distribution of the 25 results for which the reader answered at least three choices in favor of the targeted ending.

Table 3: Subset of Results (at least 3 votes)

	Escape Ending	Revenge Ending
Escape Group	12	3
Revenge Group	5	5

$p < 0.128$

In this case the *p*-value is not significant enough to reject the null hypothesis. However, the odds ratio for this table is 3.76, indicating that the chances of success are increasing as the number of shared indices increases. We believe that with a larger sample size, this table would eventually become significant as well. As there was only one participant who made all *four* choices in favor of the target ending, we did not have a sufficient sample size to evaluate the fourth table.

In addition to the comprehension questions, we also asked each participant whether they felt that the author was trying to influence them to choose one of the two endings. 82% of the 124 responses used in our evaluation reported not feeling influenced in either direction. This indicates that we can influence readers covertly, which is important when we are trying to preserve the player's feeling of agency.

Discussion and Limitations

Our results demonstrate that plan-based interactive narrative systems can utilize Indexter in combination with careful choice construction to successfully influence player choices, without impeding their sense of agency and without alerting them to any manipulation. We have shown that readers significantly prefer future events that will cause more of their past choices to become salient by sharing one or more EISM indices with them. At this point it is worth pausing to consider *why* the audience may be acting this way. It may be that readers are in a sense behaving as co-authors, and that by choosing events that are consistent with their past choices, they may be attempting to form a maximally coherent plot.

This relates to another important limitation: that the phenomena we are observing may pertain only to *endings*. It may be that readers choose final events that are consistent with their past choices in order to "tie things up" at the end of the story, but may not necessarily do the same for intermediate choices. In future work we intend to investigate whether endings are indeed a special case, and if so, whether and how we can operationalize readers' preferences for non-ending choices.

Despite its limitations, the success of this study is particularly interesting because of the simplicity of the pairwise event salience model itself, which contains several features that can be improved by future research. First, it assumes all indices are weighted equally, which is likely untrue. Second, it does not consider *how many* indices are shared by the two events—it merely considers whether or not they share *at least one*. Finally, it considers only the *most recently narrated event* as the trigger that causes past events to become more salient. It is possible that a more accurate model could be achieved, for example, by considering *all* past events, weighted according to their recency. We intend to investigate these potential improvements in the near future.

Acknowledgements

This research was supported by NSF award IIS-1464127.

References

Bae, B.-C., and Young, R. M. 2014. A computational model of narrative generation for surprise arousal. *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games* 6(2):131–143.

Bal, M. 1997. *Narratology: introduction to the theory of narrative*. University of Toronto Press.

Cardona-Rivera, R. E., and Young, R. M. 2014. A knowledge representation that models memory in narrative comprehension. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence - Student Abstracts Track*, 3098–3099.

Cardona-Rivera, R. E.; Cassell, B. A.; Ware, S. G.; and Young, R. M. 2012. Indexter: a computational model of the event-indexing situation model for characterizing narratives. In *Proceedings of the 3rd Workshop on Computational Models of Narrative*, 34–43. (awarded Best Student Paper on a Cognitive Science Topic).

Cardona-Rivera, R. E.; Robertson, J.; Ware, S. G.; Harrison, B.; Roberts, D. L.; and Young, R. M. 2014. Foreseeing meaningful choices. In *Proceedings of the 10th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 9–15.

Cheong, Y.-G., and Young, R. M. 2008. Narrative generation for suspense: modeling and evaluation. In *1st International Conference on Interactive Digital Storytelling*, 144–155.

El-Nasr, M. S.; Vasilakos, A. V.; Rao, C.; and Zupko, J. A. 2009. Dynamic intelligent lighting for directing visual attention in interactive 3-d scenes. *IEEE Trans. Comput. Intellig. and AI in Games* 1(2):145–153.

Farrell, R., and Ware, S. G. 2016. Predicting user choices in interactive narratives using indexter's pairwise event salience hypothesis. In *Proceedings of the 9th International Conference on Interactive Digital Storytelling*, 147–155.

Fikes, R. E., and Nilsson, N. J. 1972. STRIPS: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2(3):189–208.

Fleiss, J.; Levin, B.; and Paik, M. 2013. *Statistical methods for rates and proportions*. Wiley Series in Probability and Statistics. Wiley.

Friedhoff, J. 2013. Untangling twine: A platform study. In *Proceedings of the 2013 DiGRA International Conference*.

Jhala, A., and Young, R. M. 2010. Cinematic visual discourse: representation, generation, and evaluation. *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games* 2(2):69–81.

Kives, C.; Ware, S. G.; and Baker, L. J. 2015. Evaluating the pairwise event salience hypothesis in Indexter. In *Proceedings of the 11th AAAI International Conference on Artificial Intelligence and Interactive Digital Entertainment*, 30–36.

Magliano, J. P.; Miller, J.; and Zwaan, R. A. 2001. Indexing space and time in film understanding. *Applied Cognitive Psychology* 15(5):533–545.

Porteous, J.; Charles, F.; Smith, C.; Cavazza, M.; Mouw, J.; and van den Broek, P. 2017. Using virtual narratives to explore children's story understanding. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 773–781.

Riedl, M. O., and Young, R. M. 2010. Narrative planning: balancing plot and character. *Journal of Artificial Intelligence Research* 39(1):217–268.

Roberts, D. L., and Isbell, C. L. 2014. Lessons on using computationally generated influence for shaping narrative experiences. *IEEE Transactions on Computational Intelligence and AI in Games* 6(2):188–202.

Trabasso, T., and Sperry, L. L. 1985. Causal relatedness and importance of story events. *Journal of Memory and language* 24(5):595–611.

Trabasso, T., and Van Den Broek, P. 1985. Causal thinking and the representation of narrative events. *Journal of memory and language* 24(5):612–630.

Young, R. M., and Cardona-Rivera, R. E. 2011. Approaching a player model of game story comprehension through affordance in interactive narrative. In *Proceedings of the 4th Workshop on Intelligent Narrative Technologies*, 123–130.

Young, R. M.; Ware, S. G.; Cassell, B. A.; and 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(1-2):41–64.

Zacks, J. M.; Speer, N. K.; and Reynolds, J. R. 2009. Segmentation in reading and film comprehension. *Journal of Experimental Psychology: General* 138(2):307.

Zwaan, R. A., and Radvansky, G. A. 1998. Situation models in language comprehension and memory. *Psychological Bulletin* 123(2):162.