

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

Rachelyn Farrell^(✉) and Stephen G. Ware

Narrative Intelligence Lab, University of New Orleans,
New Orleans, LA 70148, USA
{rfarrell,sgware}@uno.edu

Abstract. *Indexter* is a plan-based model of narrative that incorporates cognitive scientific theories about the salience 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 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: *Indexter* · Computational models of narrative · Salience · Planning

1 Introduction

Skilled narrative authors pay close attention to how a story's events and discourse affect the audience's experience. The ability to predict what choices a user would make in an interactive narrative is useful in the context of both interactive and non-interactive narrative generation, because it can provide insight into what the user desires or expects from the future of the story. Narrative generation systems can leverage this insight to better control the audience's experience, for example by reasoning about discourse phenomena such as suspense and surprise.

The *salience* of a narrative event is defined as the ease with which the audience can recall that event. Prior research into event salience resulted in the *Indexter* model [3], 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. Events in an *Indexter* plan can share up to five narrative situation indices with one another: *protagonist*, *time*, *space*, *causality*, and *intentionality*.

A previous study [11] confirmed the *pairwise event salience hypothesis*, which 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 some past event takes place in the same room as the most recent event in the story, that past event is easier to remember than it would have been had it taken place in a different room. We now apply this theory about the salience of past events to reason about the audience’s expectation of the future. In this study 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 total number of *Indexer* indices with previous events in the story—that is, the future event which will make the past the most salient.

Participants read an interactive narrative and were prompted to choose between two possible endings. The number of indices that each ending shared with previous events in the story depended on prior choices made by the user while reading the story. Fisher’s exact test rejected the null hypothesis that the users’ choices were independent of the *Indexer* indices of past events with $p < 0.002$.

2 Related Work

Related research has focused on influencing users in interactive narratives to make choices that are in line with the author’s goal, using concepts from social psychology, discourse analysis, and natural language generation [14]. 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 [7]. Although we do not attempt to influence users’ choices in this study, the implications of our findings are relevant to that area of research.

Indexer is a computational cognitive model that reasons about how the audience comprehends a story’s *discourse*, or how the story is told [2]. The story is divided into a series of discrete events, and at each moment *Indexer* measures the salience of each past event. Plan-based models have been applied to other discourse phenomena, such as suspense [6], surprise [1], and cinematic composition [10]. Numerous plan-based models have been used to reason about story structure and to control interactive stories (see survey by Young et al. [19]). As with these other models of discourse, *Indexer* can inform story generation as well as discourse generation.

Indexer defines a plan data structure augmented with a cognitive scientific model of narrative comprehension called the event-indexing situation model, or EISM [21]. 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 [21] 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* (what enabled or impelled), and *intentionality* (why).

Indexer has also been used to predict agency in interactive stories [4]. 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 [11] suggest that *Indexer* 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 [18] call a *narrative affordance*. Recent work by these researchers [5] 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.

3 The *Indexer* Model

Indexer 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 *Indexer* maps EISM indices to plan structures, see the description by Cardona-Rivera et al. [3].

A plan is a sequence of events that achieves some goal [15]. Each event has preconditions which must be true immediately before it is executed and effects which modify the world state. The kinds of events that can occur are represented by abstract, parameterized templates called operators, as described by the STRIPS formalism [8]. For example, the domain might define an operator *attack* (*?attacker*, *?victim*, *?place*, *?time*). 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. The preconditions might be that the attacker and victim are both alive, that both are in the same place at the same time, that the attacker is armed, and that the victim is unarmed. The effects might be that the victim is no longer alive. An *Indexer* event is a fully ground instance of such an operator.

A narrative plan [13] reasons about two kinds of goals: the author’s final goal for the story and each individual character’s goals. Each event template specifies zero, one, or many of its parameters as being the consenting characters responsible for taking that action. For the *attack* example, the *?attacker* is the sole consenting character, because it carries out the action. While the *?victim* may be a character involved in the action, it need not be a consenting character.

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

Each event in an *Indexer* plan must also specify two additional required parameters: the time frame in which it occurs and the location at which it

¹ Here we use the one protagonist per event (as opposed to one per story) definition discussed by Cardona-Rivera et al. [3].

occurs. For example, the *attack* action might specify that $?time = day2$ and $?place = gym$.

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 [12, 20] 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. *Indexer* 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 *Indexer* is based is the ability to reason about causal relationships between events. While cognitive scientists have studied several forms of causality [16, 17, 21], 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 s is the causal parent of t , and that an event's causal ancestors are those events in the transitive closure of this relationship.*

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. *Let c be a character and g some goal that character c intends to achieve. Let s be an event with effect g for which c is a consenting character. Two events share the intentionality index iff both events have c as a consenting character and both are causal ancestors of s . (Note: s may be one of the events.)*

In other words, two events share *intentionality* when both are taken by the same character for the same purpose.

4 Experimental Design

We designed an interactive story wherein the user must choose between two possible endings. We hypothesize that they will choose the ending whose *Indexer* event shares more indices with previous events in the story. We allowed the user to make four intermediate choices throughout the story, each of which determines the symbol for a specific index of a single event. There are two possible symbols

for each index tested, thus a total of 16 possible story configurations. (We chose not to include the *causality* index due to the added complexity of including a choice which toggles arbitrarily between two events, where one is causally related to the ending and the other is not.) When the reader reaches the final choice, the number of indices that each possible ending event shares with the rest of the story is determined solely by their four intermediate choices.

The story is about two prisoners who are threatened by the prison bully, and each comes up with a different plan in response. Ernest plans to break out of prison and escape onto the highway, while Roy plans to get revenge by killing 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 must choose which character gets to accomplish his goal in the end.

The following is a description of how we manipulate each *Indexter* index before arriving at our 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; 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.

Space: In the same scene, the user chooses between two punishments—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: 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. For the *time* index we deviate slightly from our pattern. Since both of the endings will 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. Instead we hypothesize that the user will choose the event which shares the most *other* indices with the more salient of the two theft scenes: the one that took place on Day 2. In other words, if the most recent theft event was Roy stealing the knife for the goal of revenge, we believe the user

is more likely to choose Roy’s revenge ending over Ernest’s escape 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 which both characters will 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 to the highway. Finally, the guard catches up to them and we prompt the user for the final choice.

If the escape ending is chosen, the final event will have the parameters (*character=Ernest, location=highway, day=2, goal=escape*). If the revenge ending is chosen, it will have the parameters (*character=Roy, location=gym, day=2, goal=revenge*). The hypothesis is that the user will choose the ending event 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.

5 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 are not attempting to influence readers to choose one path or the other, many users made exactly two choices in favor of the *Escape* ending and two in favor 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 [9]. 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 (most likely due to the moral-ity 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.

6 Conclusion

We have demonstrated that interactive narrative systems can make use of *Indexter* indices to predict user choices. As the *pairwise event salience hypothesis* states, a past event is more salient if it shares at least one index with the most recently narrated event. We have shown that when users are presented with choices for future events, they generally prefer those which have more indices in common with past events. In this study we did not attempt to influence users to make specific choices, but our results suggest that future work could accomplish this using a similar method of manipulating *Indexter* indices of story events. In addition, we believe that plan-based narrative systems can utilize this information about the audience’s desires and expectations to reason about discourse phenomena such as suspense and surprise.

References

1. Bae, B.C., Young, R.M.: A computational model of narrative generation for surprise arousal. *IEEE Trans. Comput. Intell. Artif. Intell. Games* **6**(2), 131–143 (2014)
2. Bal, M.: *Narratology: introduction to the theory of narrative*. University of Toronto Press (1997). <http://books.google.com/books?isbn=1442692227>
3. Cardona-Rivera, R.E., Cassell, B.A., Ware, S.G., Young, R.M.: Indexter: a computational model of the event-indexing situation model for characterizing narratives. In: *Proceedings of the 3rd Workshop on Computational Models of Narrative*, pp. 34–43 (2012). (Awarded Best Student Paper on a Cognitive Science Topic)
4. Cardona-Rivera, R.E., Robertson, J., Ware, S.G., Harrison, B., Roberts, D.L., Young, R.M.: Foreseeing meaningful choices. In: *Proceedings of the 10th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pp. 9–15 (2014)
5. Cardona-Rivera, R.E., Young, R.M.: A knowledge representation that models memory in narrative comprehension. In: *Proceedings of the 28th AAAI Conference on Artificial Intelligence - Student Abstracts Track*, pp. 3098–3099 (2014). <https://liquidnarrative.csc.ncsu.edu/wp-content/uploads/sites/15/2015/11/cardona-rivera2014knowledge.pdf>
6. Cheong, Y.-G., Young, R.M.: Narrative generation for suspense: modeling and evaluation. In: Spierling, U., Szilas, N. (eds.) *ICIDS 2008*. LNCS, vol. 5334, pp. 144–155. Springer, Heidelberg (2008). doi:10.1007/978-3-540-89454-4_21
7. El-Nasr, M.S., Vasilakos, A.V., Rao, C., Zupko, J.A.: Dynamic intelligent lighting for directing visual attention in interactive 3-D scenes. *IEEE Trans. Comput. Intell. AI Games* **1**(2), 145–153 (2009). <http://dx.doi.org/10.1109/TCIAIG.2009.2024532>
8. Fikes, R.E., Nilsson, N.J.: STRIPS: a new approach to the application of theorem proving to problem solving. *Artif. Intell.* **2**(3), 189–208 (1972)
9. Fleiss, J., Levin, B., Paik, M.: *Statistical Methods for Rates and Proportions*. Wiley Series in Probability and Statistics, Wiley (2013). <http://books.google.co.in/books?id=9Vef07a8GeAC>
10. Jhala, A., Young, R.M.: Cinematic visual discourse: representation, generation, and evaluation. *IEEE Trans. Comput. Intell. Artif. Intell. Games* **2**(2), 69–81 (2010)
11. Kives, C., Ware, S.G., Baker, L.J.: Evaluating the pairwise event salience hypothesis in Indexter. In: *Proceedings of the 11th AAAI International Conference on Artificial Intelligence and Interactive Digital Entertainment*, pp. 30–36 (2015)
12. Magliano, J.P., Miller, J., Zwaan, R.A.: Indexing space and time in film understanding. *Appl. Cogn. Psychol.* **15**(5), 533–545 (2001)
13. Riedl, M.O., Young, R.M.: Narrative planning: balancing plot and character. *J. Artif. Intell. Res.* **39**(1), 217–268 (2010)
14. Roberts, D.L., Isbell, C.L.: Lessons on using computationally generated influence for shaping narrative experiences. *IEEE Trans. Comput. Intell. AI Games* **6**(2), 188–202 (2014)
15. Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*, 3rd edn. Prentice Hall, Upper Saddle River (2010)
16. Trabasso, T., Sperry, L.L.: Causal relatedness and importance of story events. *J. Mem. Lang.* **24**(5), 595–611 (1985)
17. Trabasso, T., Van Den Broek, P.: Causal thinking and the representation of narrative events. *J. Mem. Lang.* **24**(5), 612–630 (1985)

18. Young, R.M., Cardona-Rivera, R.E.: Approaching a player model of game story comprehension through affordance in interactive narrative. In: Proceedings of the 4th Workshop on Intelligent Narrative Technologies, pp. 123–130 (2011)
19. Young, R.M., Ware, S.G., Cassell, B.A., Robertson, J.: 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 Formal Comput. Models Narrative* **37**(1–2), 41–64 (2013)
20. Zacks, J.M., Speer, N.K., Reynolds, J.R.: Segmentation in reading and film comprehension. *J. Exp. Psychol. Gen.* **138**(2), 307 (2009)
21. Zwaan, R.A., Radvansky, G.A.: Situation models in language comprehension and memory. *Psychol. Bull.* **123**(2), 162 (1998)