A Computational Model of Inferencing in Narrative

James Niehaus and **R. Michael Young**
Department of Computer Science
North Carolina State University

Abstract

Recent work in the area of interactive narrative has sought to develop systems that automatically produce experiences for a user that are understood as stories. Much of this work, however, has focused on the structural aspects of narrative rather than the process of narrative comprehension undertaken by users. Motivated by approaches in natural language discourse generation where explicit models of a reader's mental state are often used to select and organize content in multi-sentential text, the work described here seeks to build an explicit model of a reader's inferencing process when reading (or participating in) a narrative. In this paper, we present a method for generating causal and intentional inferences, in the form of sequences of events, from a narrative discourse. We define a representation for the discourse, the sequence of discourse content, and show how it may be translated to a story representation, the reader's plan. We define cognitive criteria of necessitated inferences with regards to these representations, and show how a partial order planner can determine which inferences are enabled. The inference generation is motivated by findings in cognitive studies of discourse processing, and we provide support for their online generation by readers in a pilot study.

Introduction

Narrative comprehension is a fundamental ability possessed by nearly all human beings. Narrative is widely used as entertainment in novels and film, as an information source in news stories, as socialization in written and oral fables and histories, and as a pedagogical tool in many practical guides and textbooks. Despite its ubiquity, we know relatively little about the cognitive tools and means by which readers understand narrative. One of these tools is the construction of inferences.

Inferencing is the process by which readers add information to a narrative. Narratives rarely contain all of the information needed to understand their meaning. Instead, authors rely on the readers to use narrative conventions and reasoning to complete their understanding. Consider the example narrative in figure 1. Without inferencing, this narrative would seem to be a random collection of facts and events. Using inferencing though, one might add the information that Bogie formed a plan to steal the bacon, and Dave fell victim to it. With these inferences, the facts and events presented follow a smoother logic, and they seem much more cohesive. As in this case, inferences engage the reader in the problem solving task of understanding the narrative.

1. There was a whole, blue vase in the living room.
2. Dave was eating some bacon in the kitchen.
3. Bogie was Dave's dog.
4. Bogie wanted some bacon.
5. Bogie went to the living room.
6. There was a loud sound in the living room.
7. Bogie went to the kitchen.
8. Dave went to the living room.
9. Dave saw a shattered vase.
10. Dave went to the kitchen.
11. The bacon was gone.

Figure 1: An example of a narrative that may prompt inferences.

Cognitive psychologists have identified multiple levels at which a narrative may be mentally represented (Graesser, Olde, and Klettke 2002), and hence multiple levels at which inferences may be constructed. Guessing what action a character will perform next is an inference at the level of the situation model, the mental representation for the sum total of the story world situation. Determining the antecedent of a pronoun is an inference at the level of the surface code and textbase, the text and lexical meaning that embodies a written narrative. Similarly, studies in both narratology and computational narrative generation have often divided narrative into the *fabula* or story, the events and facts as they are in the story world, and the *szjuet* or discourse, the telling of the events and facts (Young 2006). The story is represented mentally at the level of the situation model, and the discourse is represented mentally at the levels of surface code and textbase. In this work, we generate inferences about the story at the level of the situation model.

The remainder of this paper explains our model of causal and intentional inferencing. Inferences are most likely to be made when they are *necessitated* and *enabled* (Myers, Shinjo, and Duffy 1987) (Zwaan and Radvansky 1998). Inferences are necessitated when they are required for the understanding of the narrative, and they are enabled when they are not too difficult to construct. We model the construction of a story from a discourse and define criteria for necessitated causal and intentional inferences. We use a resource bounded partial order planner to construct inferences and determine when they are enabled. Finally, we conduct a pilot study to test our predictions of inference construction.

Related Work

The IPOCL planner is a partial order narrative planner that attempts to satisfy rules of character intentionality (Riedl and Young 2004). IPOCL maintains frames of commitment to track the relationship between character actions and character intentions in order to increase character believability. We employ the IPOCL plan structure for our notions of character intentions.

A number of early computational systems have constructed various types of inferences for narrative understanding. Among these, Wilensky's PAM (1976) employs predefined "plans" to infer the connecting motivations behind character actions, but does not use a planner to construct new possible plans, nor does it consider how a reader is most likely to perform the same task. Norvig's FAUSTUS (1987) uses marker passing across a concept net to generate connecting inferences for sentences, and Hobbs et al. TACITUS (1988) uses abduction to find the same with weightings on the assumptions to show preference. More recently, Mueller (2007) uses event calculus and model fitting to construct all possible interpretations of the story, though does not include a cognitive model to justify which inferences are likely. While these systems attempt to ellicit tacit knowledge from narratives for natural language understanding, the system presented in this paper generates story level inferences to simulate aspects of the reader's mental state.

The work in psychological models of narrative discourse comprehension centers around three major theorectical frameworks: Memory Reasonance models, Embodied Cognition models, and Constructivist models (McNamara 2007). Each model makes predictions for what type of inferences are generated and under what situations they are likely to occur. Graesser, Millis, and Zwaan (1997) identify many types of inferences which may occur: "goals and plans that motivate characters' actions, character traits, characters' knowledge and beliefs, character emotions, causes of events, the consequences of events and actions, properties of objects, spatial contexts, spatial relationships among entities, the global theme or point of the text, the referents of nouns and pronouns, the attitudes of the writer, and the appropriate emotional reaction of the reader." The work in this paper is concerned with generating inferences about possible sequences of events either to satisfy the world state or to predict character actions.

McKoon and Ratcliff (1992) present a Memory Resonance view of inferencing in narrative. They contend that readers usually make inferences that contribute to only the local coherence of text, and do not attempt to infer for global coherence. They use experiments measuring reading time and word recognition response time to validate their claims. The pilot study in the present paper is of a similar design. Embodied Cognition models like that of Zwaan et al. (2002) predict activation of perceptual and sensory properties of objects and events in the narrative.

Graesser, Singer, and Trabasso (1994) present a Constructivist theory of discourse processing; the primary principle of which is that the reader attempts a *search (or effort) after meaning*. This principle is broken into three critical assumptions. *The reader goal assumption* states that "the reader constructs a meaning representation that addresses the reader's goals." *The coherence assumption* states that "The reader attempts to construct a meaning representation that is coherent at both the local and global levels". *The explanation assumption* states that "The reader attempts to explain why actions, events, and states are mentioned in the text." (Graesser, Singer, and Trabasso 1994) Graesser, Singer, and Trabasso identify a number of classes of inferences that may be constructed by the reader, and provide support for each. The work presented in the present paper enacts these assumptions in the reader's search after meaning.

Problem Representation

Data Structures

In order to effectively reason about inference construction, we define a logical representation for both the story and discourse of a narrative. For the story level, the IPOCL plan representation (Riedl and Young 2004) maintains information about causal relationships between actions as well as information about characters' intentions. This information will enable the system to reason about causality and intentionality at the story level. For the discourse, we define a *Sequence of Discourse Content*, a total ordering over elements which may be found in an IPOCL plan.

Definition, IPOCL Plan *is a tuple,* $< S, B, O, L, C >$*, where* S *is a set of steps,* B *is a set of binding constraints on the free variables in* S*,* O *is the set of ordering constraints on steps in* S*,* L *is a set of causal links between steps in* S*, and* C *is a set of frames of commitment.* (Riedl and Young 2004)

Steps, binding constraints, ordering constraints, and causal links are defined in the usual way, as in UCPOP (Penberthy and Weld 1992). The frame of commitment, IPOCL's representation of intentionality, separates IPOCL plans from UCPOP's partial order plans. The frames serve to map steps to character intentions. If a step is part of a frame of commitment, it is performed by the character in service of the frame's goal (a single predicate). For example, in the narrative in figure 1, Dave may have formed the (unstated) intention to find the source of the loud sound, and his entering the living room may be in service of this goal. The step MOVE-TO(DAVE, KITCHEN, LIVING-ROOM) is part of the frame with the goal KNOWS-SOURCE-OF(DAVE, LOUD-SOUND1).

We define the discourse level representation to exhibit two observations we make about the narrative reading process.

1. *Readers are presented with a subset of the information in the story level.* This observation is what makes inferencing possible in and often necessary for narrative understanding.
2. *The information is presented in a linear order.* The ordering of information is important to determine which inferences might be constructed at any point in the story.

An IPOCL plan, though appropriate for the story level representation, does not fit well with these observations. A plan is partially ordered, and it contains all of the story information. Instead, we find it more natural to view the reading process as the construction of a story plan from a discourse representation, the sequence of discourse content.

Definition (Sequence of Discourse Content) *A sequence of discourse content is a tuple, $< R, S, B, O, L, C, T >$, where S is a set of steps, R is a set of predicates from the preconditions and effects of the steps in S, B is a set of bindings on the free variables in S, O is a set of ordering constraints on steps in S, L is a set of causal links between steps in S, C is a set of intentional frames of commitment, and T is a total ordering over $S \bigcup R \bigcup C$.*

The sequence of discourse content contains three types of elements: steps, predicates, and frames of commitment. The steps have the same structure as plan steps, and they represent events in the narrative such as "Jim kicked the ball". The predicates represent single facts about the world such as "The vase rested on the table" and may also serve to introduce and identify objects in the world: "There was a girl named Sally". The frames of commitment identify intentions of the characters as in "Bob wanted to eat some ice cream.". The total ordering over these elements asserts the order in which they should be presented to the reader. Figure 2 shows a visualization of an example sequence of discourse content.

Narrative Understanding

A reader presented with a discourse attempts to reconstruct the elements of the story. Just as there are potentially infinite stories which can be transcribed as a specific discourse (e.g. repeatedly add a single event to the story, unmentioned in the discourse), there are a potentially infinite IPOCL plans for each sequence of discourse content. We assume that the sequence, SD, contains only true information about the story plan, Q, namely that $S_{SD} \subset S_Q$, $B_{SD} \subset B_Q$, $O_{SD} \subset O_Q$, $L_{SD} \subset L_Q$, $C_{SD} \subset C_Q$. Note that SD may contain all of the elements of the Q and still prompt inferences, and that the SD may prompt inferences that are not found in Q. For reasons of parsimony and limited cognitive resources on the part of the readers, we choose to attempt to construct IPOCL plans that minimize the information not found directly in the discourse. Sequences of events which transform the story world state are the only form of inferences we treat in this paper.

The algorithm in Figure 1 depicts the general process for incorporating elements from the discourse into the story and then generating inferences. The sequence of discourse content, Q is processed element by element according to the total ordering over elements. At each step the new element e_i is directly incorporated into the reader's representation of the story, an IPOCL plan called the *reader's plan*, RP. Then the possible inferences I are generated (though not incorporated into the reader's plan) for this point in the discourse, and the algorithm recurses. When done processing the elements in the sequence of discourse content, the algorithm returns the reader's plan and the collection of inferences made during processing.

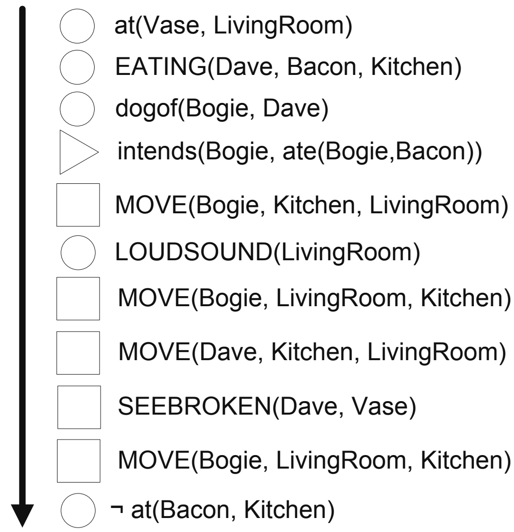

Figure 2: A sequence of discourse content. Read top to bottom, the sequence is a representation of the narrative in figure 1. The circles depict predicates, the triangle is an intention, and the squares are plan steps.

Algorithm 1 Algorithm for reconstructing the story. Q is the sequence of discourse content in queue form, RP is the reader's plan, and I is the set of inferences. The Inference function may return a set with no inferences, one inference, or multiple competing inferences. All of these would be included in the final result.

Understand(Q, RP, I):
if Empty(Q) **then**
  RETURN $[RP, I]$
else
  $e_i \leftarrow Pop(Q)$
  $RP \leftarrow Incorporate(e_i, RP)$
  $I \leftarrow I \bigcup Inference(RP)$
  RETURN $Understand(Q, RP, I)$
end if

Incrementally constructing the reader's plan RP from a sequence of discourse content is relatively straight forward. The *Incorporate* function returns a new IPOCL plan. Here we make the simplifying assumption that the order of elements in the sequence of discourse content is consistent with a total ordering of the story; the story is told without flash-forwards or flashbacks. This assumption is reflective of the

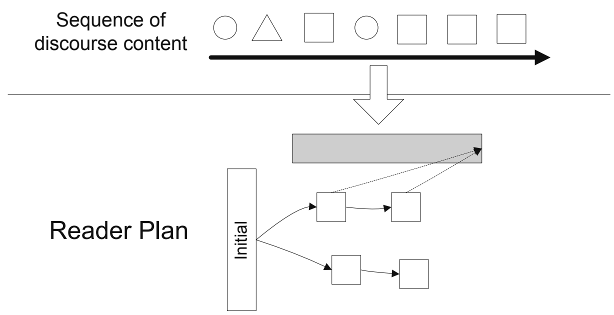

Figure 3: Translating a sequence of discourse content to a reader plan.

ordering of elements assumed by the logical progression of time in a narrative, and it creates a more constrained reader's plan for inferencing. Figure 3 shows the translation from sequence to plan. Incorporate includes the new element as follows:

- If e_i is a step, add e_i and all of its bindings and causal links to its preconditions to RP. Add an ordering from each step proceeding e_i to e_i. Add any new bindings established by causal links.
- If e_i is a character intention, then create a new frame of commitment f_i in RP. Add an ordering from each step proceeding f_i to f_i (as per the order preserving constraint).
- If e_i is a predicate, create a dummy step s_i with the sole precondition and effect of e_i. Add s_i to RP as above.

Inferencing

In this section, we define criteria for determining when to construct two types of inferences, causal and intentional, and provide an algorithm for generating each. After a new element has been incorporated into the reader's plan the plan may yet appear incomplete. We generate causal inferences when the truth value of predicates in the reader's plan change without explanation. We generate intentional inferences when characters have unfulfilled intentions. The inferences are generated by the use of a partial order planner.

Determining when to construct inferences as well as what information to construct inferences about is a significant problem. Inferences could be constructed at any time during reading, and they may be about anything in the story world. Consider the discourse in figure 5. The reader may choose to construct inferences concerning the color of the front door, the number of keys Billy had, the reasons for Billy entering the door, the reasons the front door was locked, or that Billy unlocked the front door using his keys.

Myers et al. suggest that unprompted inferences are more often made when inferencing is *necessitated* to understand the story and *enabled*, so that the inference is not difficult to construct (Myers, Shinjo, and Duffy 1987) (Zwaan and Radvansky 1998). Graesser and Clark (Graesser and Clark 1985) note that readers rarely make unconstrained predictive inferences and that when they are asked to, the inferences often prove to be untrue. For the restricted view of inferencing in this paper, we define criteria for determining when causal and intentional inferences are necessitated and enabled.

Causal inferences are necessitated when the truth value of predicates changes without an intervening step. Consider figures 4 and 5. In figure 4, the reader does not know whether the door was locked at the beginning of the story or not. The truth value of LOCKED(FRONTDOOR) is asserted to be false by the opening of the door, but never changed. In figure 5, LOCKED(FRONTDOOR) is asserted to be true as the first line in the discourse, but then the character enters the front door. Entering the door requires the door be unlocked, and the truth value of LOCKED(FRONTDOOR) is changed to false. At this point a causal inference is necessitated to determine how the door became unlocked. The reader might infer that Billy unlocked the door using the keys.

Intentional inferences are necessitated when characters have unfulfilled intentions. In figure 6, the reader learns of Adams intention to get something to eat. At this point, the intentional inference of how Adam will fulfill his intention of getting something to eat is necessitated (though perhaps not enabled).

1. Billy grabbed his keys.
2. Billy entered the front door.

Figure 4: Example narrative without causal necessity.

Both causal and intentional inferences are enabled when the reader has enough information to create the inference and the reasoning process is short enough to be plausible. In the inferencing presented in this paper, the truth value of predicates in the story is never assumed or inferred. Instead, sequences of events which transform the story world state are the only form of inferences. In order for these inferences to be made, the discourse must have presented all the requisite preconditions for the sequence of events. In addition to the possibility of making the inference in a context, the inference must also be plausible. Readers do not often make large, unconnected inferences, or inferences that are only one of a large number of possibilities (Graesser, Singer, and Trabasso 1994). Hence the other half of the enablement criteria is that the search space for inferences be relatively small. Enablement is defined precisely in the algorithms below.

Causal inferencing occurs when the reader reasons about

1. The front door was locked.
2. Billy grabbed his keys.
3. Billy entered the front door.

Figure 5: Example narrative with causal necessity.

how events may occur or may have occurred to produce the current narrative situation. To simulate causal inferencing, we treat the current state of the reader's plan as a planning problem. Given a sufficient set of planning operators to describe the possibilities of events in the narrative, the planner can provide the causal reasoning necessary to determine which steps have occurred in the intervening period.

Figure 7 displays the process for generating causal inferences. The inferences are necessitated because they only concern predicates which have changed truth values. The inferences are enabled because the resource bounded planning prevents the search from becoming too deep or wide (depending on the type of bound employed).

An intervening step, as used in causal inferencing, is one that changes the truth value of a predicate in the required way, false to true or true to false, within the interval between steps where the truth value is reversed. Predicates which change truth value without an intervening step are those that prompt causal inferences. An open precondition flaw, to demarkate a precondition for which the planner is to plan, is created from a predicate that has changed its truth value. Using the open precondition as the only initial flaw in the plan guarantees the planner will only attempt to solve for this flaw and any new flaws that occur as a result of added elements. The planner will ignore other open preconditions or threatened links in the plan, focusing on the sequence which resulted in the change in p_i.

Intentional inferencing is the processes by which readers predict character actions based upon perceived character goals. To simulate intentional inferencing, we treat the current state of the reader's plan with the goal of achieving the character's intention as a planning problem. Given a sufficient set of planning operators to describe the possibilities of actions by the character, the planner can provide the causal reasoning necessary to determine which steps the character my take to achieve his goal.

Figure 8 displays the process for generating intentional inferences. The inferences are necessitated by the unfulfilled intentions. The inferences are enabled because the resource bounded planning prevents the search from becoming too deep or wide (depending on the type of bound employed). Intentional inferencing is similar to causal inferencing except that the open precondition is a character intention and only that character's actions are used in the planning.

Pilot Study

The pilot study was designed to assess the availability of inference related information at the end of reading short narratives. The design of this study is similar to studies in the discourse processing literature which also test models of inference generation (McKoon and Ratcliff 1992) (Suh and Trabasso 1993). This study tests online inferences without reading strategy instruction: inferences that are made during reading without specific instruction to do so - these may be considered automatic inferences. If the inferences are occurring, we theorized that they should activate related information in the reader's mental model of the story, and make this information more available at the end of the narrative.

The narratives consisted of 4 narratives designed to test causal inferencing, 4 narratives designed to test intentional inferencing, and 4 filler narratives. The narratives were each constructed as an IPOCL plan at the story level and a list of elements at the discourse level. A simple templating function was used to map the list of elements into a list of human readable sentences.

Each of the test narratives had two versions, one version, the prompted condition, which constructed the appropriate inference according to our algorithms and one version, the un-prompted condition, which did not, due to either lack of necessitation or enablement. Participants read through the narratives one sentence at a time in a self paced manner. Reading times were recorded for each sentence. After each narrative, the participant was asked to answer two word recognition questions, stating whether a given word was in the narrative or not. One of these tests related to the prompted inference.

In the causal inferencing test case, readers were asked about a word not in the text but related to the prompted causal inference. We hypothesized that in the prompted con-

1. Adam was in the kitchen.
2. Adam wanted something to eat.
3. There was an apple on the kitchen counter.

Figure 6: Example narrative with intentional necessity.

1. Let L be the list of predicates that have changed their truth value without an intervening step (all of which are preconditions). Let p_i be the ith element of L, and s_i be the step with precondition p_i.
2. Conduct resource bounded planning with the open precondition p_i as the only initial flaw in the plan. Expand the tree fully up to the resource bounds.
3. The possible inferences for p_i are the new sets of elements in the un-flawed plans within this search space.

Figure 7: Generating causal inferences.

1. Let F be the list of frames of commitment with unfulfilled goals. Let f_i be the the ith element of F, g_i be the goal of f_i and c_i be the character of f_i.
2. Conduct resource bounded planning with the open precondition g_i as the only initial flaw in the plan. Use only actions with the character c_i bound as the actor. Expand the tree fully up to the resource bounds.
3. The possible inferences for f_i are the new sets of elements in the un-flawed plans within this search space.

Figure 8: Generating intentional inferences.

dition readers would take longer to answer this test and make more errors. In the intentional inferencing test case, readers were asked about a word in the text in the stated intention of the character. We hypothesized that in the prompted condition readers would take less time to answer this test and make fewer errors.

Method

Materials. The narratives consisted of 4 filler narratives and 8 experimental divided between 4 narratives designed to test causal inferencing and 4 narratives designed to test intentional inferencing. The narratives were each constructed as an IPOCL plan at the story level and a list of elements at the discourse level. A simple templating function was used to map the list of elements into a list of human readable sentences.

Each of the experimental narratives had two versions, the same length in sentences and approximately the same number of words per sentence. Each experimental narrative generated the appropriate inference when the algorithms above were applied to the lists of elements in the prompted condition, and did not in the un-prompted condition. The stories ranged from 9 to 21 sentences each, and at most 4 lines were changed between versions.

Procedure. The narratives were displayed by a Java applet within a browser window. The experiment began with instructions to read for comprehension (the participants were not informed they would be timed) and a short example narrative. Participants placed their hands on a standard QWERTY keyboard so that their left index finger was on the 'z' key, their right index finger was on the '/' key and their thumbs were on the space bar. The participants advanced the narrative one sentence at a time by pressing the spacebar. After each press of the spacebar, the current sentence on screen was erased, there was a 200-ms pause, and the next sentence was displayed. Participants could not return to earlier sentences.

At the end of each narrative two word recognition tests were administered. The participant was presented with the question as to whether a particular word appeared in the narrative and indicated either 'yes' by pressing the 'z' key or 'no' by pressing the '/' key. There was a 500-ms pause between questions and between a question and the start screen for the next narrative. Participants were given untimed rest periods after the first two sets of 4 narratives, indicated by onscreen instructions.

Design and Subjects. 18 students from an introductory Artificial Intelligence class participated in the study for course credit. Each participant only read one version of each narrative. The narrative versions were counter-balanced in a rotating sequence, and the order of the narratives given to a participant was randomized.

Results

The overall results are presented in tables 1 and 2 (these results are not statistically significant). Reading times per word are included for completeness. We did not form any hypotheses for the reading times.

Table 1: Intentional Narrative Results

	RT mean	RT StdDev	% error	rd mean
prompted	3,304 ms	2,148 ms	11.1	394 ms
unprompted	3,179 ms	2,126 ms	9.7	401 ms

Note RT = response time, rd = reading time

Table 2: Causal Narrative Results

	RT mean	RT StdDev	% error	rd mean
prompted	4,384 ms	3,149 ms	6.9	596 ms
unprompted	4,054 ms	3,544 ms	4.2	591 ms

Note RT = response time, rd = reading time

Table 3: Intentional Narratives, Slow Readers

	RT mean	RT StdDev
prompted	4,484 ms	2,901 ms
unprompted	5,763 ms	4,324 ms

Table 4: Intentional Narratives, Fast Readers

	RT mean	RT StdDev
prompted	4,285* ms	3,462 ms
unprompted	2,346* ms	990 ms

* $p < .069, F(1, 18) = 3.53$

In the intentional narratives, participants were slightly faster to respond in the unprompted condition, and the error rate was slightly less in this condition as well. They were able to recall that a word relating to the intention was part of a narrative slightly faster and more accurately when the events in the narrative did not heavily involve that intention or allow for inferencing involving that intention. This result is counter to our original hypothesis.

In the causal narratives, participants were slightly faster to respond in the unprompted condition, and the error rate was slightly less. They were able to recall correctly that a word was *not* in the narrative slight faster and more accurately when the narrative did not prompt an inference related to the word. The prompting of the inference appears to slow down response time and increase error rate. This result is in line with our original hypothesis (though without statistical significance).

Two potentially interesting trends emerge from the data. The first is that reading times correlate significantly with response times, $R^2 = .60, p < .001$. The second is that the variances for the response times are quite large; one standard deviation away from the mean ranges from 1 to 5 seconds in the intentional narratives. With these in mind we split the participants into slow (above median) and fast (equal or below median) readers based on their average reading times and examined the effect of the conditions on these groups.

Tables 3 and 4 show the results of the second analysis for the intentional narratives. In the slow readers group, prompting decreased the response time. Adding actions related to the intention seemed to improve availability, as per our hypothesis. In the fast readers group, prompting increased the response time dramatically. Adding actions related to the intention seemed to hinder availability. This last result approached significance at $p < .069$.

Discussion

In the causal inferencing narratives, the overall results provide weak support for our hypothesis. The addition of causal inferences relating to a word not in the text seem to slow response times and increase error rates. The information may be more salient in the reader's mind and the reader may have a more difficult time discerning whether the information was inferenced or included.

In the intentional inferencing narratives, the effect of the inferences seems to depend on whether the reading times are slow or fast. In the slow reader group, the addition of the possibility of intentional inferences seems to make the information more available to the reader. In this case, we posit that the reader may be using the semantics of the narrative to recall the information, and is thus aided by the inferences. In the fast reader group, the prompting of intentional inferences significantly slows the response time. In this case, we posit that the reader is relying more on the surface text of the narrative. Reading time is faster because the reader does not have to take the time to encode the semantics of the narrative, but response time is hindered by related information. The reader takes longer to recall a specific item out of a collection of highly related items than it does to recall a specific item out of a collection of unrelated items.

Further experimentation is needed to verify these claims. In post-hoc interviews, several of the participants reported that they thought that many of the stories were 'tricking' them by indicating something had happened without actually stating the event (prompting an inference) and then asking about the event in the questions. This realization seemed to interrupt their normal reading strategy. The sole use of word recognition tests may have cause the participants to focus heavily on the lexical properties of the text (memorizing words) rather than understanding the content of the narrative, which the experiment was designed to test. We would consider including comprehension questions to balance this effect.

Conclusions

In this paper, we have presented a method for generating causal and intentional inferences from a narrative discourse. We have defined a representation for the discourse, the sequence of discourse content, and show how it may be translated to a story representation, the reader's plan. We defined the cognitive criteria of necessitated inferences with regards to these representations, and show how a partial order planner can determine which inferences are enabled. The inference generation is motivated by findings in cognitive studies of discourse processing, and we provide support for their online generation by readers in a pilot study.

Though the pilot study results are promising, they only concern the online construction of inferences, when perhaps some of the inferences we generate are more likely in offline reasoning. Future studies may address offline inferences. Also, the enablement criteria as defined above address the difficulty of constructing an inference when the elements are readily available to the reader, but does not take into account the focus of the reader. Readers will have a more difficult time constructing inferences between disparate elements in a narrative than between those that are cognitively, if not textually, close. A model of narrative focus is needed to address this issue.

The model of inference generation in this paper may be used for narrative generation as well as understanding. A system may be parameterized with inferences, and then attempt to prompt the reader to make these inferences by examining the potential discourses with this model. Inference generation has been related to recall and comprehension (Graesser, Olde, and Klettke 2002), factors which may increase the teaching ability of such narratives.

References

Graesser, A., and Clark, L. 1985. *Structures and procedures of implicit knowledge*. Norwood, NJ: Ablex.

Graesser, A.; Olde, B.; and Klettke, B. 2002. How does the mind construct and represent stories? In Green, M. C.; Strange, J. J.; and Brock, T. C., eds., *Narrative Impact : Social and Cognitive Foundations*. L. Erlbaum Associates.

Graesser, A.; Singer, M.; and Trabasso, T. 1994. Constructing inferences during narrative text comprehension. *Psychological Review* 101(3):371–395.

Graesser, A.; Singer, M.; and Trabasso, T. 1997. Discourse comprehension. *Psychological Review* 48:163–189.

Hobbs, J. R.; Stickel, M.; Martin, P.; and Edwards, D. 1988. Interpretation as abduction. In *Proceedings of the 26th annual meeting on Association for Computational Linguistics*, 95–103. Morristown, NJ, USA: Association for Computational Linguistics.

McKoon, G., and Ratcliff, R. 1992. Inference during reading. *Psychological Review* 99(3):440–466.

McNamara, D. S. 2007. *Reading Comprehension Strategies: Theories, Interventions, and Technologies*. Routledge.

Mueller, E. T. 2007. Understanding goal-based stories through model finding and planning. In Magerko, B., and Riedl, M., eds., *Intelligent Narrative Technologies: Papers from the AAAI Fall Symposium*, 95–101. Menlo Park, CA: AAAI Press.

Myers, J. L.; Shinjo, M.; and Duffy, S. A. 1987. Degrees of causal relatedness and memory. *Journal of Memory and Language* 26:453–465.

Norvig, P. 1987. Inference in text understanding. In *Proceedings of the National Conference on Artificial Intelligence*, 561–565.

Penberthy, J., and Weld, D. 1992. Ucpop: a sound, complete, partial order planner for adl. In *Proceedings Third International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, 103–114.

Riedl, M. O., and Young, R. M. 2004. An intent-driven planner for multi-agent story generation. *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems, New York*.

Suh, S., and Trabasso, T. 1993. Inferences during reading: Converging evidence from discourse analysis, talk-aloud protocols, and recognition priming. *Journal of Memory and Language* 32:279–300.

Wilensky, R. 1976. Using plans to understand natural language. In *ACM 76: Proceedings of the annual conference*, 46–50. New York, NY, USA: ACM Press.

Young, R. M. 2006. Story and discourse: A bipartite model of narrative generation in virtual worlds. *Interaction Studies*.

Zwaan, R. A., and Radvansky, G. A. 1998. Situation models in language comprehension and memory. *Psychological Bulletin*.

Zwaan, R. A.; Stanfield, R. A.; and Yaxley, R. H. 2002. Language comprehenders mentally represent the shapes of objects. *Psychological Science* 13(2):168–171.