# Using Narrative Functions as a Heuristic for Relevance in Story Understanding

Emmett Tomai
University of Texas - Pan American
1201 W. University Dr.
Edinburg, TX, 78539, USA
tomaie@utpa.edu

Ken Forbus
Northwestern University
633 Clark Street
Evanston, IL 60208, USA
forbus@northwestern.edu

## ABSTRACT
Story understanding requires a degree of knowledge and expressiveness beyond the current state of natural language understanding. We present an approach that addresses these needs, using a large-scale knowledge base, simplified English grammar and a combination of compositional frame semantics and abductive reasoning. This in turn raises a significant challenge disambiguating complex semantic structures, which requires a pragmatics of narrative for constraint and guidance. We present a theory of narrative functions that serve as a heuristic for relevance in narrative, and provide evidence that this heuristic is effective for disambiguation that leads to consistent understanding.

## Categories and Subject Descriptors
I.2.7 [**Artificial Intelligence**]: Natural Language Processing – *discourse, language parsing and understanding, text analysis.*

## General Terms
Algorithms, Design, Experimentation.

## Keywords
Story, Narrative, Natural Language Understanding, Semantics, Abductive Reasoning, Knowledge Representation.

## 1. INTRODUCTION
This paper discusses an approach to narrative understanding. The narratives in question are short, text stories such as folktales and fables that relate a series of events undertaken by a set of actors. In this account, understanding a narrative is generating a formal logical representation that supports inference. In previous work [24], we have presented a practical approach to language understanding that supports deep, task-oriented understanding of narrative scenarios taken from the psychological literature. One of the key elements of this approach is that it is grounded in a large-scale knowledge base using a highly expressive representation language. This allows us to generate logical forms that cover the *semantic breadth* encountered in narrative. Semantic breadth is the range of distinct world models that can be presented in the natural language text. This is not only a matter of the unconstrained topics that might arise, but also the challenge of representing modal situations such as beliefs, descriptions and hypotheticals. These constructs add considerable complexity to logical representations, but are ubiquitous in narrative and trivially handled by human language competence. We have demonstrated that our approach, combining large-scale knowledge, compositional frame semantics and abductive reasoning with pragmatic constraint, is capable of building these representations and effectively performing semantic disambiguation. One significant trade-off for this semantic breadth is that our approach uses a simplified syntax to make the parsing challenge tractable. We feel that this is a complementary approach to the numerous investigations demonstrating broad syntactic breadth over a relatively small range of possible semantic forms.

In this paper we apply our approach to the problem of general narrative understanding, focusing on fables. We first discuss the range of semantic forms we are dealing with, and the challenges of disambiguation in this knowledge-rich environment. This motivates the need for a pragmatics of narrative to provide constraint and guidance for abductive interpretation. We argue that this pragmatics can be cast as a search for *relevance*, and we present a theory of *narrative functions* which we hypothesize can serve as a heuristic measure of relevance in narrative. Finally, we present a small-scale evaluation of this hypothesis providing evidence that this heuristic 1) is able to effectively guide disambiguation and 2) leads to a sufficient understanding of narrative to perform a sentence ordering task.

## 2. PRACTICAL LANGUAGE UNDERSTANDING
To support general commonsense inference, we have grounded our representations in a large-scale knowledge base. It is made up of the contents of ResearchCyc[1] [17] plus our own extensions, together around 2 million facts at present. We use this with the highly expressive CycL language to maximize the range of distinct world models that can be represented. The mapping from language to concepts in the Cyc ontology begins with the extensive *denotations* and *subcategorization frames* in the KB. Denotations map directly from lexical terms to concepts within the Cyc ontology. Subcategorization frames follow Fillmore's theory of frame semantics [12] and map from a term to a logical form. Each frame selects for syntactic arguments that fill in

[1] http://research.cyc.com/


INT3 2010, June 18, Monterey, CA, USA

semantic roles in its logical form. These forms must then be composed by a parsing process that handles role assignment and quantification. Figure 1 contains the semantic translation for a frame for the term "went" containing ACTION and SUBJECT roles. It also shows an example of the frame once it is filled and quantified.

```
(and
 (isa :ACTION Movement-TranslationEvent)
 (primaryObjectMoving :ACTION :SUBJECT))

(thereExists (TheList he1 move1)
 (and
  (isa move1 Movement-TranslationEvent)
  (primaryObjectMoving move1 he1)))
```

**Figure 1. semantic frame for "went".**

Figure 2 contains a more complex semantic translation for a frame for the term "wanted" which expects a SUBJECT and an infinitive complement clause (INF-COMP). It also shows a possible filled and quantifed form for the phrase "he wanted to move" where the clausal substitution results in a higher-order nested quantitfication.

```
(desires :SUBJECT :INF-COMP)

(thereExists he1
 (desires he1
  (thereExists move1
   (and
    (isa move1 Movement-TranslationEvent)
    (primaryObjectMoving move1 he1)))))
```

**Figure 2. semantic frame for "wanted".**

The possibility of higher-order nesting for facts results in highly expressive logical forms with considerable semantic breadth. However, it also multiplies the points of explicit ambiguity that must be dealt with. A given term may have many frames corresponding to different cases, and those frames may involve different nesting of the semantics of other terms. Additional complexity comes from quantifier scoping and anaphora resolution ambiguities. We use Allen's bottom-up chart parser [1] with the COMLEX lexicon [18] to do lambda-calculus composition of the semantics of the constituent terms. We made only minor modifications to the parser to allow it to retrieve lexical knowledge and frames from the KB. This setup allows us to support the complex compositions as described above, but comes with a heavy cost in computational complexity. To make the problem tractable we use a simplified English grammar to constrain syntactic ambiguity. Since we are focused on extending semantic breadth, supporting more than one surface form for a given internal representation is a secondary concern.

Following [3], we combine this compositional approach with a transformation process using dynamic logic principles from Discourse Representation Theory (DRT) [15]. This process takes the logical form composed by the parser and constructs a model-theoretic description of sentence content. Numerical and logical quantifiers, negation and implication are handled according to DRT by constructing nested discourse representation structures (DRS). The same representation is extended for possible worlds indicated by the modal operators *possible* and *willBe*. Figure 3 contains the DRS corresponding to the example semantic form for "he wanted to move" in Figure 2.

```
Universe: he1

(desires (DrsCaseFn DRS-1))

DRS-1:

  Universe: move1

  (isa move1 Movement-TranslationEvent)
  (primaryObjectMoving move1 he1)
```

**Figure 3. DRS for "He wanted to move".**

The result of this sentence-level processing is not a single DRS. There are potential ambiguities in the syntax of the parse trees, quantifier scoping, semantic role assignments and frame selection. These ambiguities are reified into explicit *choice sets* to be disambiguated by the discourse-level interpretation process. The output of the sentence-level process is these choice sets combined with a set of axioms that entail both facts and DRS structure based on choices within the sets. This representation is suitable for abductive proof, enabling back-chaining from a certain fact in a certain DRS to the choice set selections that would make it true if assumed. In previous work [24], we have shown that abductive proof of task-specific queries is an effective way to disambiguate these choice sets. We now turn to the challenge of disambiguation in the general task of narrative interpretation.

## 3. COHERENCE AND RELEVANCE

Several researchers have explored task-general abductive interpretation as proving the logical form of a sentence on the basis of prior knowledge [6], [20], [13]. Hobbs extended this framework to show how additional pragmatic constraints can be layered on top of these proofs. He used a set of *coherence relations* between successive utterances. These relations, such as explanation and elaboration, indicate how one utterance coherently follows from another, and can be expressed as axioms within an abductive reasoning system. This shifts the interpretation task from proving the form of the utterances up to proving that they form a coherent whole. The proof of each coherence relation depends on proving the form of the utterances it relates, thus achieving the utterance interpretations with additional constraint. Hobbs also argues that this framework is not limited to coherence within the story. Pragmatic theories about explaining speaker intentions can be likewise axiomatized and proven [13]. This is a very elegant formulation, but the notion of coherence requires a global proof over all the utterances in a discourse. Heuristics for incremental understanding are not obvious, since the disambiguating factor for one utterance may not come for an arbitrary number of utterances.

Relevance theory [28] addresses these issues in part, by defining a measure of relevance and two principles that state that 1) "human cognition is geared to the maximization of relevance" and 2) "utterances create expectations of optimal relevance". By this account, the problem of finding the most likely interpretation of an utterance is solved by maximizing the relevance of that utterance at the time it is heard. This is defined as the number of *positive cognitive effects* generated by a particular interpretation. Examples of positive cognitive effects are: a conclusion that can be drawn from the utterance and the context together but from neither alone, as well as evidence for or against prior assumptions.

Judgments of relevance are typically discussed in situated dialogue where the pragmatic concerns of the speaker and hearer may be invoked. If one is waiting for a train, then statements about the arrival time of that particular train are notably relevant and interpretation of ambiguous elements can be guided in that direction. However, within a narrative this notion of positive cognitive effect is insufficient to gauge the relevance of an utterance. The opening sentence in a story, for example "An ant went to a river to drink.", does not present any true conclusions nor does it reference any known entities or update any existing model. Instead, it establishes expectation. The imagined fact of an ant, being situated by a river and desiring a drink, will lead to further developments which the hearer can reasonably expect to be relevant. We propose that it is these expectations that act as heuristics for relevance in narrative understanding.

# 4. NARRATIVE FUNCTIONS AS RELEVANCE

Structuralist narratology is concerned with the dualism of *story* and *discourse:* the events being described versus the means used to describe them (cf. [7]). In this view, every narrative utterance serves one or more functions to express content (cf. [4]). These functions form a pragmatics of narrative: a set of intentional moves made by the narrator and expected by the listener.

The realization of a narrative function in our system is represented as a *fulfillment relation* between a *presentation event (PE)* and a set of entities from the story. The PE is a reified, deliberately underspecified part of a sentence that fulfills one or more narrative functions. For example, the presentation event PE1 would be asserted to introduce the actor ant1234 by (1).

(1) (introducesActor PE1 ant1234)

Our system uses expectation of these fulfillment relations as a heuristic for relevance. For each new sentence, it attempts to abductively prove that the sentence contains PEs that fulfill narrative functions in the context of the ongoing discourse. In the example sentence, "An ant went to a river to drink", the abductive proof is able to conclude that there is a PE within this sentence that introduces a new actor (the ant) and thus that it fulfills a narrative function. This conclusion is based on the axiom (simplified by removing quantification):

(2) (isa ?entity Agent-Generic)
$\wedge \neg$ (isa ?entity Organization)
$\wedge$ (resolveReference ?entity ?entity ?sentence-id)
$\rightarrow$ (introducesActor ?PE ?entity)

This reasoning relies on specialization relations in ResearchCyc (the collection *Ant* is a specialization of *Agent-Generic*, but not *Organization*) as well as reasoning to determine that the discourse variable ?entity in the new sentence can only be resolved to itself. The reference resolution in this case is based on the use of an indefinite reference (*an* ant) in the first sentence of the story. Because it resolves to itself only, it necessarily does not resolve to any existing (already introduced) character. Further details on the reference resolution reasoning are omitted here.

The abductive proof of (1) indicates that the relation could hold true, but only if certain assumptions are made. Those assumptions correspond to syntactic and semantic choices in the choice sets generated by the sentence-level compositional semantics. Thus the heuristic of fulfilling a narrative function (introducing an actor in this case) provides evidence for certain disambiguation choices.

## 4.1 Expectation Functions

Unlike more rigid story grammars, this theory of narrative functions uses a flexible system of setting and meeting expectations to provide higher level structure. This is expressed as two higher-order narrative functions. The fulfillment relations for these functions take other fulfillment relations as their arguments.

(3) (setsExpectation ?frel-sets ?expected-frel)

(4) (meetsExpectation ?frel-meets ?frel-sets)

The abductive proof of (3) indicates that the fulfillment relation ?frel-sets sets the expectation that a later PE will fulfill the relation ?expected-frel. For example, (5) indicates that when PE1 presents a goal to achieve state1, being held by actor actor1, it sets the expectation that a later PE will present actions performed by actor1 with the purpose of achieving state1.

(5) (setsExpectation
(presentsGoal PE1 actor1 (Goal-AchieveFn state1))
(presentsAction ?PE actor1 ?action
(Goal-AchieveFn state1))

It is important to note that "later" here refers to discourse time, not story time. Expectations are necessarily established first, then met in the telling. However, that expectation fulfillment may be a revelation of something earlier in the story timeline. In this study we do not deal with flashbacks or reveals of this type.

To complete the expectation fulfillment, the abductive proof of (4) indicates that a certain relation ?frel-meets serves to meet the expectation established by ?frel-sets. For a certain ?frel-sets, ?frel-meets in (4) is necessarily equivalent to ?expected-frel in (3) with all free variables bound to ground terms.

In most cases, a sentence can be interpreted in more than one way. Each interpretation is based on certain disambiguation choices which in turn are chosen to support certain fulfillment relations. In the cases where those choices contradict each other, the non-contradictory subset of choices that result in the highest total *relevance score* are chosen. This relevance score is based on the number of narrative functions the sentence is interpreted to perform, and how well those functions fulfill prior expectations. The relevance score is calculated for a sentence by summing the scores of all its fulfillment relations. The score for a fulfillment relation *frel* is calculated according to the recursive algorithm shown in Figure 4, using a constant C. The constant influences how much weight is given to meeting expectations, and was set to a single, unchanging value, 3, for all calculations in all test sets presented here.

```
relevance( frel ) =
 gather all frel' such that:
   (meetsExpectation frel frel')
 if frel' = {} then
   C
 else
   C * ∏ relevance(frel')
```

**Figure 4. recursive algorithm for calculating relevance.**

Because the meetsExpectation relations are unidirectional from earlier to later arguments, there are no cycles to be concerned with.

The remainder of our narrative functions can be categorized according to three structural elements of fiction widely recognized in narrative theories (cf. [7]). Here we adopt the terms *action*, *character* and *setting*.

## 4.2 Action

The realm of action centers on four types of *goals*, each a unary function on a DRS expression of a model. An achievement goal is satisfied when the model becomes valid in the world of the story. A maintenance goal must remain valid. An avoidance goal must not become valid and a cessation goal must cease to be valid. The first action function therefore is presenting a goal held by one or more actors. The fulfillment relation is shown in (6), where ?goal binds to a functional goal term such as (Goal-MaintainFn model1).

(6) (presentsGoal ?PE ?actor ?goal).

When a goal is presented, it creates the expectation that actions will be performed with respect to that goal, as in (7), and that in turn sets the expectation that the outcome of the goal will be revealed, as in (8).

(7) (setsExpectation
    (presentsGoal ?PE1 ?actor ?goal)
    (presentsAction ?PE2 ?actor ?action ?goal))

(8) (setsExpectation
    (presentsAction ?PE1 ?actor ?action ?goal)
    (presentsOutcome ?PE2 ?actor ?action ?outcome))

It is important to note that expecting these things does not necessitate their taking place, but if they do it will result in a higher relevance score. The ?outcome variable in (8) is a relationship (e.g. *enables*, *achieves*, *fails*) between the action and either another action, a goal, an opportunity or a threat. This goal-action-outcome pattern is similar to the *causal network* representation [26], including the possibility of sub-goals. The fulfillment relationship for presenting sub-goals is given in (9).

(9) (presentsSubGoal ?PE ?goal-sub ?goal-super)

Intentional, goal-directed actions are complemented by opportunistic and threatening circumstances. These are also granted relevance with respect to an actor who holds a goal – the opportunity to succeed or the threat of failure. These fulfillment relations are given in (10) and (11).

(10) (presentsOpportunity ?PE ?actor ?situation ?goal)

(11) (presentsThreat ?PE ?actor ? situation ?goal)

Both threats and opportunities motivate action distinctly from (although often in concert with) goal motivated behavior. Thus they raise the additional expectation that a notable response is forthcoming by one of the actors. The fulfillment relation is given in (12). They also raise an expectation regarding their outcome, is the same way as deliberate actions shown in (8).

(12) (setsExpectation
    (presentsThreat ?PE1 ?actor ?situation ?goal)
    (presentsResponse ?PE2 ?actor ?action ?situation))

In addition to the action functions surrounding goals, the awareness of an actor to some situation or event is notable. This is captured simply as the fulfillment relation (13).

(13) (presentsAwareness ?PE ?actor ?situation)

## 4.3 Character

The first function in the realm of character, as discussed above, is introducing new actors in the story. Beyond this, characterization is the function indicating that an actor possesses a certain internal trait. This characterization can be diegetic (explicitly told by the narrator) or mimetic (indicated within the story itself). Figure 5 shows the categories of mimetic characterization (cf. [25]). The fulfillment relation is given in (14).

Explicit (direct)
- Expository self-description
- Expository description of another

Implicit (indirect)
- External appearance
- Non-verbal behavior
- Verbal behavior
  - Content of utterance
  - Form of utterance

**Figure 5. categories of mimetic characterization.**

(14) (characterize ?PE ?category ?subcat ?actor ?trait)

The category and subcat variables are bound to constants indicating the means of characterization. Character traits are more-or-less persistent qualities such as being patient or compassionate rather than temporary emotions such as happiness or frustration. In like fashion, relationships can be characterized according to established roles (e.g. hero, nemesis, love-interest, etc).

(15) (characterize-Reln-Role ?PE
    ?category ?subcat ?actor ?actee ?role)

This is a unidirectional relationship indicating that the actor is presented as being in the specified role with respect to the actee (it does not imply or contradict the reverse relationship).

## 4.4 Setting

The functions of setting serve to situate the action in space and time. As a measure of relevance, statements that introduce set pieces and props are of particular interest because they raise the expectation that those entities will be used in a notable way; a river might be fallen into or something purchased in a town. These fulfillment relations are given in (16) and (17).

(16) (setsExpectation
(introducesSetting ?PE1 ?setting)
(presentsUse ?PE2 ?setting ?usage))

(17) (setsExpectation
(introducesProp ?PE1 ?prop)
(presentsUse ?PE2 ?prop ?usage))

The usage variable is a relation between the setting or prop and an event. Typical relations include *instrument*, *enables* and *into*.

## 4.5 Example

Figure 6 contains the simplified English rendering of Aesop's fable "The Ass, the Fox and the Lion". Figure 7 shows the structure of narrative functions (in the action dimension) identified in the automatic interpretation of this story. Events in the story are listed across the bottom in chronological discourse order. The system represents discourse sequence as a simple ordered sequence, while temporal relations in the story are represented using Allen's temporal interval calculus [2], as implemented in ResearchCyc. Story events are not assumed to follow one after another, but may overlap and even come out of order. In this study, however, the latter does not occur. The narrative functions associated with the presentation of each event are depicted in vertical columns above the events. The solid arrows indicate expectations set and met, while the dashed line indicates a sub-goal relation. The abductive disambiguation process identifies each function as part of a possible interpretation along with the set of choice set choices that justify it. The functions (and choices) selected are those that resulted in the highest relevance score for each sentence. Because this is a heuristic method, there is no guarantee that this is a global maximum over the entire story.

S1: An Ass and a Fox, having entered into a partnership for protection, went into the forest to hunt.
S2: They had proceeded a short distance when they met a Lion.
S3: The Fox, seeing danger, approached the Lion and promised to capture the Ass for him.
S4: In return, the Lion would promise to not harm the Fox.
S5: The Fox, having assured the Ass that he would not be injured, led him to a deep pit.
S6: He then caused the Ass to fall into it.
S7: The Lion, seeing that the Ass was trapped, immediately attacked the Fox.
S8: He then attacked the Ass at his leisure.

**Figure 6. simplified text for "The Ass, the Fox and the Lion".**

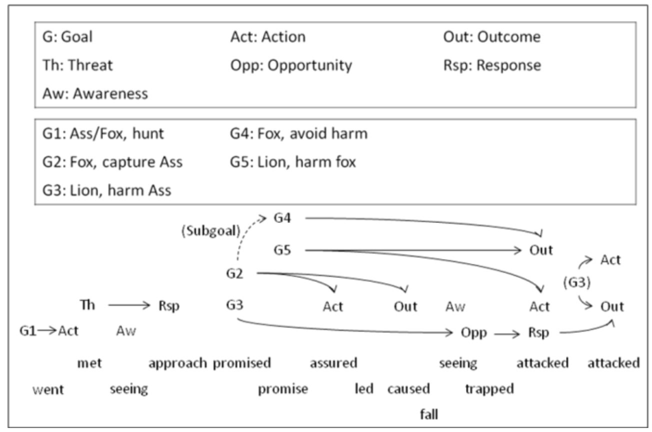


**Figure 7. narrative functions identified in "The Ass, the Fox and the Lion".**

# 5. Evaluation

We have performed a preliminary evaluation of the effectiveness of this approach using two Aesop's Fables. We tested the hypothesis that this theory of narrative functions as relevance can guide abductive disambiguation towards a coherent and relevant understanding of a story. The two fables consist of 15 sentences in our simplified English, ranging from 5 to 18 words each. Figure 6 contains the simplified text for "The Ass, the Fox and the Lion" and figure 8 contains the simplified text for "The Dove and the Ant".

S1: An ant went to a river to drink.
S2: The fell into the river and was carried along in the stream.
S3: A dove pitied her condition and threw a small bough into the river.
S4: The ant used the bough to reach the shore.
S5: Afterward, the ant saw a man aiming a gun at the Dove.
S6: The ant stung him in the foot, causing him to miss.
S7: This saved the dove's life.

**Figure 8. simplified text for "The Dove and the Ant".**

Table 1 and Table 2 show the number of explicit ambiguities generated by the sentence-level compositional interpretation of the two stories. The number of semantic frame choice sets in each sentence (corresponding to the number of terms with more than one frame) is followed in parenthesis by the average number of choices per set.

**Table 1. ambiguities in "The Ass, the Fox and the Lion."**

| | Parse trees | Semantic frames | Anaphora |
|---|---|---|---|
| S1 | 2 | 4(5) | 0 |
| S2 | 1 | 4(4.75) | 2 |
| S3 | 1 | 4(3.25) | 2 |
| S4 | 1 | 2(4) | 2 |
| S5 | 4 | 3(7.67) | 2 |
| S6 | 1 | 5(3.5) | 3 |
| S7 | 1 | 1(8) | 2 |

**Table 2. ambiguities in "The Dove and the Ant".**

| | Parse trees | Semantic frames | Anaphora |
|---|---|---|---|
| S1 | 1 | 6(5.25) | 0 |
| S2 | 1 | 2(3) | 2 |
| S3 | 2 | 8(3) | 4 |
| S4 | 1 | 4(3) | 2 |
| S5 | 1 | 7(3.5) | 4 |
| S6 | 1 | 3(3.67) | 3 |
| S7 | 1 | 5(4) | 3 |
| S8 | 1 | 3(4) | 3 |

The system was given sufficient knowledge, both from the ResearchCyc KB and added for this evaluation, to resolve these ambiguities by abductive proof of relevance. Our additions consisted of filling in missing subcategorization frames, typically simple extensions from frames that were already present in the KB, and adding rules relevant to the particular situations in the stories. Even in a knowledge base the size of ResearchCyc, reasoning about, for example, the motivations behind throwing a branch into a river is very hit-and-miss. The disambiguation results in a complex, logical representation of each story, suitable for reasoning tasks. In other work [9], we have addressed the usefulness of representations generated by this system for other, more task-focused sets of stories. To evaluate the effectiveness of this particular theory of narrative functions as a heuristic for relevance, we tested the system on a sentence ordering task.

In the ordering task, the system is given a story and one sentence that has been removed from that story (the test sentence). It must then identify the original position of the test sentence. The system approaches this task by placing the test sentence in each of the n possible positions. Each of the resulting n versions of the story is interpreted by the system, and the relevance scores of all its sentences are summed. If this theory of narrative functions is a reasonable heuristic for relevance, then the correct ordering should have a higher relevance score than any incorrect ordering. Thus, the system selects the version with the highest relevance score as the correct ordering. In plain terms, we expect that if a sentence is moved to an incorrect position, the events and observations in that sentence will be less relevant to what is happening than if the sentence appears in its proper place. Likewise, events and observations in other sentences that rely on the moved sentence will be less relevant than in the correct ordering.

Figure 9 shows the relevance scores for the each possible test sentence (grouped along the horizontal axis) in each possible position for "The Ass, the Fox and the Lion". For sentence S1, the highest score is obtained in position 1. For sentence S2, the highest score is obtained in position 2, and so on. Of course those highest scores are all identical, since the correct ordering is always the original ordering. Figure 10 shows the same data for "The Dove and the Ant".

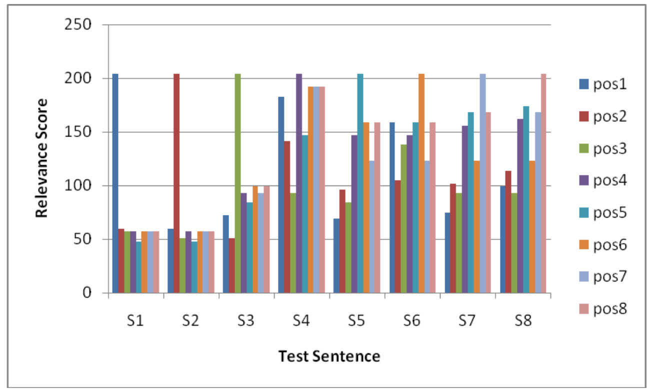


**Figure 9. relevance scores for orderings of "The Ass, the Fox and the Lion".**

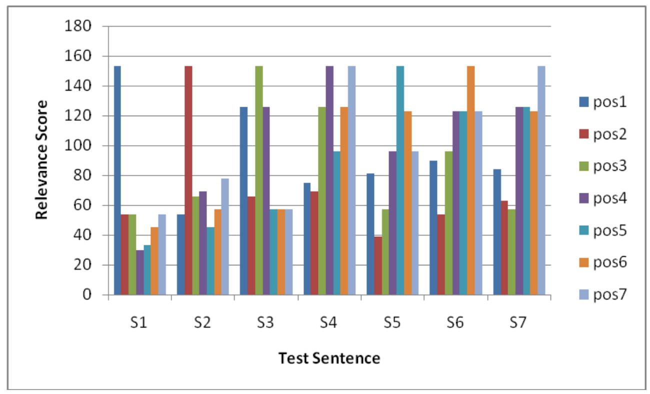


**Figure 10. relevance scores for orderings of "The Dove and the Ant".**

In all 8 ordering tests for "The Ass, the Fox and the Lion", the test sentence scores higher in its correct position than any other. In 6 of the 7 ordering tests for "The Dove and the Ant", the test sentence scores higher in its correct position than any other. For test sentence S4 of "The Dove and the Ant", one incorrect position receives a score equal to the correct position. This is when the sentence "The ant used the bough to reach the shore." is moved to the end of the story. It appears that the system treats the second half of the story, involving the Dove's peril, as an aside. After the Dove has been saved, the Ant's situation is resolved. Clearly this fails to recognize that the Ant cannot be drowning in the river and saving the Dove at the same time.

These results provide evidence that this theory of narrative functions can be used as an effective heuristic for relevance in story understanding. The flexible structure of establishing and fulfilling expectations is able to identify those cases where each subsequent action is most relevant to those that came before.

## 6. RELATED WORK

Narrative interpretation was long studied by Schank and his students as a problem of applying world knowledge. They hypothesized that understanding a new story was a matter of invoking previous experiences stored as various types of patterns in memory, then using those patterns to direct subsequent inferences. The FRUMP system [10] used *scripts* to represent typical scenarios that could then be recognized and summarized in news stories. Wilensky's PAM [27] used patterns of causal and intentional behavior to understand actor motivations in utterance pairs. This work demonstrated the effectiveness of knowledge regarding typical occurrences, but suffered from scaling problems due to its reliance on stereotypical situations and the lack of availability of broad world knowledge.

Several models of narrative structure have been proposed, but few have been implemented and tested in a computational framework. Propp's morphology of Russian folktales [21] was seminal in structuralist narratology but generally recognized to be limited in scope to that particular genre of folktale. Mandler and Johnson's story grammar [19], based on prior work by Rumelhart [22], and Trabasso's causal networks [26] propose structures of goal-based action and outcome that we have drawn from here. However, the global structure they present is overly rigid and does not attempt to account for opportunistic motivations, awareness or the dimensions of character and setting. Lehnert's plot units [16], and Ryan's extension for character points of view [23], also provide useful insight into positive and negative outcomes and their back-and-forth relationship with mental motivations. Finally, Dyer's BORIS [11] did a full implementation of in-depth story understanding directed by a pragmatic theory of *thematic abstraction units*. These structures represented proverbs as plan-failure cases which, if already known, could serve to guide the interpretation of a story communicating that proverb. We have attempted to lay the groundwork for a compositional theory of narrative that can perform similarly without requiring prior knowledge of the point the story is trying to make.

The Boxer [5] system, with the C&C Tools parser [8], also creates DRT-style representations with concern for quantification and semantic role filling. However, it does not ground its predicates in an ontology (or otherwise axiomize them) for general reasoning. We see this as a complementary effort. These tools aim to provide robust, large-scale processing working from the bottom up and succeed in attaining an impressive amount of breadth for the depth. We have used simplified English to get past syntactic concerns and investigate from the top down. We remain interested in ways of improving syntactic breadth, but only insofar as we are able to do so without compromising semantic breadth.

## 7. CONCLUSION

We have presented an approach to narrative understanding that emphasizes semantic breadth to support broad and deep inference. One of the challenges of this approach is disambiguating complex semantic structures, requiring the application of world knowledge and pragmatics. The latter is notably difficult in the field of general-purpose narrative understanding. We have proposed a theory of narrative functions that can be used as a heuristic for relevance in interpretation. Abductive proof that a sentence in a story fulfills one or more of these functions provides an effective framework for disambiguating the selection of ambiguous parse trees, semantic frames, role assignments and pronoun and definite noun phrase anaphora. We have presented experimental results providing evidence that these narrative functions are an effective measure of the relevance of a sentence given the sentences that have come before it.

## 8. FUTURE WORK

We intend to expand our theory of narrative functions to incorporate a fourth structural element, commentary within the discourse. This will allow us to extend our analysis to the explicit morals provided with many fables and folktales. Explaining why the moral is relevant to the story proper should shed light on how the meaning of the story and the advice of the moral combine to present a more compelling message than either alone.

We are also working on a computational model of the impact of cultural folktales on decision-making strategies. We believe that relevance can also be an indicator of emphasis, highlighting the salient factors (point of view, behavior, circumstance) that are presented as responsible for positive or negative outcomes. By automatically identifying these factors in culturally important stories we can create a knowledge base of scenarios which offer advice on decisions. We have already begun testing the use of such scenarios against psychological findings in cultural decision-making.

Finally, we are very interested in the use of narrative communication in interactive simulations, such as training scenarios and games. In this line of work, we are primarily concerned with expanding our understanding of the types of narrative functions and pragmatic narrator strategies to fulfill them. How does a narrator set up expectations? How do they lead the audience towards a certain interpretation that can then be reinforced or reversed? The space of narrative freedom, and the act of guiding an audience through it, are ill-understood in formal, computational terms. We will continue to work towards better models that combine reasonable computational characteristics with powerful expressiveness.

## 9. ACKNOWLEDGMENTS

This research was supported by the Air Force Office of Scientific Research.

## 10. REFERENCES

[1] Allen, J. F. 1994. *Natural Language Understanding*. (2nd ed). Redwood City, CA.: Benjamin/Cummings.

[2] Allen, J. and Ferguson, G. 1997. Actions and events in interval temporal logic. *Spatial and Temporal Reasoning.* Springer.

[3] Asher, N. and Lascarides, A.. 2003. *Logics of Conversation*. Cambridge University Press, Cambridge, UK.

[4] Barthes, R. 1977. Introduction to the structural analysis of narratives. *Image, music, text; essays selected and translated by Stephen Heath*. Hill and Wang, New York.

[5] Bos, J, Clark, S, Steedman, M, Curran, J and Hockenmaier, J. 2004. Wide-Coverage Semantic Representations from a CCG Parser. Proceedings of the 20th International Conference on Computational Linguistics (COLING-04).

[6] Charniak, E. and Goldman, R.P. 1989. A semantics for probabilistic quantifier-free first-order languages, with particular application to story understanding. IJCAI.

[7] Chatman, S. 1978. *Story and Discourse: Narrative Structure in Fiction and Film*. Cornell University Press.

[8] Clark, S. and Curran, J.R. (2007): *Wide-Coverage Efficient Statistical Parsing with CCG and Log-Linear Models*. Computational Linguistics, 33(4), 2007.

[9] Dehghani, M., Tomai, E., Forbus, K., Klenk, M. 2008. An Integrated Reasoning Approach to Moral Decision-Making. In the Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI). Chicago, IL.

[10] DeJong, G.F. 1977. Skimming newspaper stories by computer, Research Report 104. Department of Computer Science, Yale University.

[11] Dyer, M. 1983. *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. MIT Press Cambridge, MA, USA.

[12] Fillmore, C. 1982. Frame semantics. In *Linguistics in the Morning Calm*. Seoul, Hanshin Publishing Co., 111-137.

[13] Hobbs, J.R. 2004. Abduction in Natural Language Understanding. In Horn, L. and Ward, G. (eds.), *Handbook of Pragmatics*, pp. 724-741, Blackwell.

[14] Hobbs, J.R. 1985. On the Coherence and Structure of Discourse. Report No: CSLI-85-37. Center for the Study of Language and Information, October, 1985

[15] Kamp, H. and Reyle, U. 1993. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language*. Kluwer Academic Dordrecht; Boston.

[16] Lehnert, W.G. 1981. Plot units and narrative summarization. *Cognitive Science*. Elsevier.

[17] Lenat, D., and R.V. Gupta. *Building Large Knowledge Based Systems*. Reading, MA: Addison Wesley, 1990.

[18] Macleod, C., Grishman, R., and Meyers, A. 1998. COMLEX Syntax Reference Manual, Version 3.0. Linguistic Data Consortium. University of Pennsylvania: Philadelphia, PA.

[19] Mandler, J.M. and Johnson, N.S. 1977. Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*, 9,111-151.

[20] Ng, H.T. and Mooney, R.J. 1990. On the Role of Coherence in Abductive Explanation. Proceedings of the Eighth National Conference on Artificial Intelligence.

[21] Propp, V. 1969. *Morphology of the Folktale*. Trans. Laurence Scott. Ed. Louis A. Wagner. 2nd edition. Univ. of Texas Press.

[22] Rummelhart, D.E. 1975. Notes on a schema for stories. In D.G. Bobrow & A. Collins (eds). *Representation and understanding: Studies in cognitive science*. New York: Academic Press.

[23] Ryan, M. 1991. *Possible Worlds, Artificial Intelligence, and Narrative Theory: A Possible Worlds/artificial Intelligence Approach to Narrative and to Fiction*. Indiana University Press.

[24] Tomai, E. and Forbus, K. (2009) EA NLU: Practical Language Understanding for Cognitive Modeling. Proceedings of the 22nd International Florida Artificial Intelligence Research Society Conference. Sanibel Island, FL.

[25] Tomashevsky, B. 1925/1965. Thematics. In *Russian Formalist Criticism: Four Essays*. Translated and with introduction by Lemon, L.T. and Reis, M.J. Lincoln.

[26] Trabasso, T., Secco, T. and van den Broek, P.W. 1984. Causal cohesion and story coherence. In H. Mandl, N.L. Stein and T. Trabasso (eds.). *Learning and comprehension of text*. 83-111. Lawrence Erlbaum Associates, Hillsdale, NJ.

[27] Wilensky, R. 1981. PAM and Micro PAM. In Roger C. Schank, & Christopher K. Riesbeck (Eds.), *Inside computer understanding*. Hillsdale, NJ: Erlbaum.

[28] Wilson, D. and Sperber, D. 2002. Relevance theory. *Handbook of Pragmatics*. L. Horn and G. Ward (eds). Oxford: Blackwell.