

Toward a Knowledge Representation for Simple Narratives

R. Raymond Lang and Robert P. Goldman

Computer Science Department, 301 Stanley Thomas Hall, Tulane University
New Orleans, Louisiana 70118, (504) 865-5840

Abstract

In this paper we report progress on the design of a knowledge representation formalism, based on Allen's temporal logic [Allen, 1984], to be used in a generative model of narratives. Our goal is to develop a model that will simultaneously generate text and meaning representations so that claims about recovery of meaning from text can be assessed. We take as our domain a class of simple stories, based on Grimm's fairy tales. We base our work on story grammars, as they are the only available framework with a declarative representation. We provide the logical foundations for developing a story grammar [Rumelhart, 1975] into a generative model of simple narratives. We have provided definitions to specify the "syntactic" categories of the story grammar and the constraints between constituents.

Introduction

It is not possible to assess the success of natural language processing programs which aim at a 'deep' (e.g., script- or plan-based) understanding of text, absent some criterion of correct understanding. To address this problem, we propose to construct a generative model of a restricted class of narratives. This model must simultaneously generate text and meaning representations so that claims about recovery of meaning from text can be assessed. Furthermore, this model must be specified in such a way that it can be understood and critiqued as a declarative representation, not just as imbedded in a program.

In this paper we report progress on the design of a knowledge representation formalism to be used in a generative model of narratives. For reasons outlined below, we take as our domain a class of simple stories, based on Grimm's fairy tales. We base the formalism on story grammars, as they are the only available framework with a declarative representation.

We wish to emphasize that we are not seeking to model text production. Our model need only capture the correct relationships between the texts and the desired interpretations; although our goal is to develop

a model which will map in both directions (i.e. from texts to meanings as well as from meanings to texts). Furthermore, the model need not be *complete* in that factors which are not important to the interpretation problem which interests us can either be left implicit or be omitted from the model.

In designing our KR formalism, we use Allen's temporal logic [Allen, 1984], a language based on first order logic. While we are conscious of criticisms of first order logic as a knowledge representation language (e.g., [McDermott, 1987, Birnbaum, 1991]), we quite simply do not know of an alternative knowledge representation language which has a semantics, an inference calculus (deductive, "scruffy" or otherwise) and is either more expressive or more convenient. We do not wish to be seen as offering an argument on either side of the debate about the use of logic in AI; we are simply choosing the best tool available to us at this time.

Domain: Simple Stories

Much as cooking is often used as a domain for research in planning, simple stories are frequently chosen as a sample domain for testing theories of narrative structure and natural language understanding. Previous work in story generation has followed one of two tracks: (1) procedural and (2) descriptive.

Computer Story Telling

Within the first track, Meehan's TALE-SPIN [1976] is the most well-developed and widely known work. It is fundamentally a simulation of a forest world, producing natural language output describing the interactions of characters pursuing goals such as eating and drinking in a context where duplicity and hostility are possible as well as honesty and friendliness. Among story-telling programs, TALE-SPIN comes closest to what we seek in terms of having access to the meanings (conceptual dependency forms, in this case) from which the natural language text is constructed. However, the model by which the meanings themselves are generated is left implicit; and the relationships among the components of a story are deeply entwined in the procedures which drive the simulation.

1. Story → Setting + Episode
⇒ ALLOW (Setting, Episode)
2. Setting → (States)*
⇒ AND (State, State, ...)
3. Episode → Event + Reaction
⇒ INITIATE (Event, Reaction)
4. Event →
{Episode | Change-of-State | Action | Event + Event}
⇒ CAUSE (Event₁, Event₂) or ALLOW (Event₁, Event₂)
5. Reaction → Internal Response + Overt Response
⇒ MOTIVATE (Internal Response, Overt Response)
6. Internal Response → {Emotion | Desire}
7. Overt Response → {Action | Attempt*}
⇒ THEN (Attempt₁, Attempt₂, ...)
8. Attempt → Plan + Application
⇒ MOTIVATE (Plan, Application)
9. Application → (Preaction)* + Action + Consequence
⇒ ALLOW (AND(Preaction, Preaction, ...),
{CAUSE | INITIATE | ALLOW}
(Action, Consequence))
10. Preaction → Subgoal + (Attempt)*
⇒ MOTIVATE [Subgoal, THEN (Attempt, ...)]
11. Consequence → {Reaction | Event}

Figure 1: Rumelhart's grammar for stories

Later work in story generation by computer addresses concerns not relevant to our project. Lebowitz's UNIVERSE [1985] is fundamentally a planner in that the output is generated by means of meeting the requirements of a goal. Lebowitz is explicitly interested in author-level goals such as "create suspense" and how these contribute to the content of a story. As such, this work would be more useful if we were trying to automate, say, literary criticism rather than simple, literal understanding. In MINSTREL [Turner and Dyer, 1985], Turner seeks to model the process of human creativity and uses King Arthur-style tales as his domain. Although we are working in a similar domain, we are not making any claims about what a human does when writing a story.

Modeling the Logical Structure of Stories

Work in the modeling of the structure of stories is founded upon seminal works such as Propp's *Morphology of the Folktale* [1968] and Polti's *The Thirty-Six Dramatic Situations* [1921]. Although Propp's perspective is that of an anthropologist and Polti's that of a literary critic, these works (particularly Propp) have been used frequently as a starting point for the development of theories of story structure and as sources for the sorts of categories that pertain to stories.

David Rumelhart [1975] presented a "story grammar" which inspired subsequent work which attempted to express the regular "syntax" of stories by means of grammars. His grammar is reproduced in Figure 1. Rumelhart's grammar, as well as many of the others like it [Bower, 1976, Johnson and Mandler, 1980, Mandler and Johnson, 1977, Stein and Glenn, 1979, Thorndyke, 1977], are context-free grammars. Some, including Rumelhart's, are augmented by "semantic" constraints. The terms *syntax* and *semantics* as used in these grammars may be misleading. The syntactic rules are constraints on event sequences and the semantic constraints restrict the relations between consecutive events.

Interest in story grammars soon lapsed, partially due to an attack on their foundations by Black and Wilensky [1979] (henceforth B&W). They make the case that a story grammar must be an unrestricted rewrite system (Chomsky type 0), but they find only one proposed story grammar [Johnson and Mandler, 1980] that meets this criteria. B&W argue that story grammars fail by trying to capture the idea of what a story is by trying to express the structure of a story text. In particular, they claim there is some purpose in relating a story; and the listener (or reader) of a story has some interest in hearing it. In other words, a story has a "point," i.e. some element that invokes the interest of a reader. The alternative theory of story points given by Wilensky [1982] views the form of a story as a function of its content. Although there is merit to this approach, the points theory fails to account for the regular structure of stories. B&W either ignored or failed to consider results of experiments in story recall supporting the hypothesis that people use story schemas to understand simple stories [Bower, 1976, Mandler and Johnson, 1977, Stein and Glenn, 1979, Thorndyke, 1977].

Part of the reason for the success of B&W's attack was the crude state of techniques for formalizing material like story grammars. In Rumelhart's grammar (as well as those based on this grammar), information about the relationship a story component has to other components is restricted to annotations accompanying the rules. In these grammars, the "syntactic" structure of a portion of a text makes a particular rule applicable, then the relationship of this component to others is gleaned from the annotation to the rule. Unfortunately, the "syntax" given in these grammars doesn't rule out many constructions; while the "semantic" annotations are not formalized rigorously enough.

Since the deficiencies of story grammars are so well documented, the reader is justified in asking why we are resurrecting a formalism laid to rest ten years ago. We believe the idea of a story grammar embodies important insights regarding what there is about a sequence of events such that reporting it constitutes a story. It is these insights we aim to formalize. Among the grammars available to us, Rumelhart's was chosen

as a starting point because his use of semantic annotations showed the most promise. Although we use this grammar as a starting point, we also seek to incorporate important intuitions from other grammars, in particular Johnson and Mandler's [1980].

Another reason for the apparent success of B&W's criticisms of story grammars is that their point theory pertains to an different class of narratives than Rumelhart's simple stories. The work of both B&W and that of Rumelhart suffers by lacking a clearly defined domain. It is not possible to say specifically what kinds of stories these authors deal with. B&W's narratives have a more contemporary tone and style than the simple folk tales Rumelhart attempts to describe. In particular, simple folk tales frequently lack antagonists in B&W's sense of an independent agent with his or her own goals in competition with those of the protagonist. Whatever "bad guys" there are in simpler stories are merely animate obstacles that the protagonist must overcome in pursuit of his or her goal. We also find stories in which the main character has no clear goal. These stories are little more than a character's reactions to a sequence of external events. Such stories fall outside B&W's domain, but not outside of ours.

In contrast, Mandler and Johnson [1977, 1980] confine their attention to simple stories limited to a single protagonist in each episode. They explicitly allow stories in which events in one episode lead to another episode in which a different character becomes the protagonist; however, we are presently restricting our work to stories with a single protagonist throughout. Typically, the protagonist in our class of stories will be pursuing an explicit goal which may be given in the setting of the tale or arise as a reaction to an initial event.

We believe the real difficulty with the story grammar approach is that much of the insight is in annotations to the rules, but that the relations used to specify these annotations is vague and fuzzy. Our task, then, is to formalize these intuitions in a rigorous manner. In the next section, we outline the logic we use as a foundation for our work. Following this, we give some examples of how we have used this logic to translate Rumelhart's grammar into a system of logical axioms.

Temporal Logic

We base our formalization upon James Allen's temporal logic [Allen, 1984]. We have chosen this logic because it was specifically designed as a formalism for reasoning about actions in the context of natural language processing, while skirting the complications of modal and higher-order logics.

We have considered some other logics, notably McDermott's temporal logic [McDermott, 1982], and the episodic logic of Schubert and Hwang [1989]. We have chosen Allen's logic for reasons of simplicity. Unlike McDermott's logic, the semantics of Allen's logic is

based on a single time line. Since we are primarily concerned with being able to represent sequences of (conceptually) past actions, we see no benefit to branching time semantics to offset the more complicated model. We have been inspired by Schubert and Hwang's work, but find their logic to be heavily biased towards natural language *understanding*, as opposed to simple representation of action sequences. For example, they introduce mechanisms for referring to existential variables *outside* their scopes, and structure the syntax of their logic so as to closely parallel syntactic analyses of natural language sentences. We do not have present use for these complexities, so prefer to skirt them.

The entities in Allen's logic are action and state types, temporal intervals and individual objects and beings. Important relations express the occurrence of events and persistence properties over temporal intervals. E.g., $HOLDS(foo(x), t)$ states that the *foo* property holds of some entity *x* over time interval *t*, and $OCCUR(bar(x), t)$ says that an event of type *bar*, involving an entity *x* occurs over time interval *t*. Other important predicates denote relations between temporal intervals, e.g., $BEFORE(t_1, t_2)$ says that time interval *t*₁ is before, and does not overlap with *t*₂. Allen's logic also attempts to give an account of planning and intentional behaviors. Modal aspects of the logic (agents' beliefs) are handled by a quotation method, so the logic as a whole is first-order. We have developed a treatment of intentional action based on Allen's logic, which is reported elsewhere [Goldman and Lang, 1992].

We extend and modify Allen's logic in 4 ways in order to make it more appropriate to the task at hand.

1. We extend the ontology of Allen's logic to include event tokens. This is done as a matter of convenience: our project here is to talk about what it is which makes a series of events a candidate for being recounted as a story. Having to talk about events as pairs of event descriptions and intervals is unnecessarily cumbersome. This extension is captured in the predicate $event(token)$, which is true when *token* denotes a particular instance of *event-class* which occurred during *time-interval*, that is,

$$\forall x event(x) \rightarrow OCCUR(eclass(x), time(x))$$

where *eclass* is a function which takes an event token and returns the class of events into which it falls, and *time* is a function over event tokens which returns the time-interval at which the event took place.

2. Allen provides a predicate, *IS-GOAL-OF* for making statements about an agent's intentions. In his logic, a statement of the form

$$IS-GOAL-OF(agent, property, gtime, t)$$

denotes the fact that the *agent*, at time *t*, desires that *property* hold, at time *gtime*.

It is not clear exactly what this statement says about the intentions of *agent*. Is he or she aware of the extent of *gtime*? The difficulty can be seen most

clearly if one wishes to interpret a proposition of the form $\exists gt$ IS-GOAL-OF (agent,property,gt,t). In this case, is there a specific time when the *agent* wants *property* to be true, which we just do not happen to know, although he does? Or does *agent* just want the action to be true at some fairly arbitrary time in the future?

In our domain we are more concerned with the latter case than the former (which arises naturally when one needs to reason about actions under time pressure).

- Allen characterizes plans as a set of decisions about performing or not performing some set of actions. Unfortunately, Allen's account of goals and plans are only tenuously related. Allen's logic does not allow us to state anything stronger than that x has a goal of g at the same time x pursues plan p . One cannot express the fact that x is pursuing plan p in order to achieve goal g .

We need this more specific relation for at least two purposes. First, we need to give a definition of plan failure in order to specify sensible action sequences. We see no way to do this absent a relation between the plan and its desired end. We also need to make statements which specify which plans are sensibly applicable to which classes of goals. Finally, for more complex stories, we will have to be able to distinguish between intended effects, and unintended side-effects.

- Finally, for convenience, we extend the syntax of the logic with the \oplus operator, written in prefix form, to indicate exclusive "one-of" disjunction, following Kautz [1986]. Also, we include quantification over restricted parts of the domain. We follow Schubert and Hwang [1989] in using the notation $\exists x[x : type]P(x)$ as a shorthand for $\exists x : type(x) \wedge P(x)$ and likewise $\forall x[x : type]P(x)$ for $\forall x : type(x) \rightarrow P(x)$

A Logic for Stories

In this section, we outline and give examples of our approach to the two immediate tasks which we face in building our logic: (1) developing semantic categories to match Rumelhart's "syntactic" elements, and (2) giving firm definitions of Rumelhart's "semantic" relations.

As mentioned earlier, we believe the semantic relationships and syntactic categories Rumelhart describes are capable of describing a non-trivial class of simple stories. Our work up to this point has been aimed at eliminating the vagueness and informality in Rumelhart's descriptions of these relationships and categories. For the semantic relationships, we clearly delineate the number and types of arguments and define them in terms of the extensions to Allen's logic outlined above. The categories have been recast as single-place predicates, indicating that the argument, which may be a token, is of the indicated type.

Categories As an example, consider the categories *action* and *change-of-state*. Rumelhart describes these informally in a way that leaves open the possibility that a token could be both simultaneously. We establish these as mutually exclusive sub-types of a restricted version of Rumelhart's category, *event*. Rumelhart characterizes an *action* as "an activity engaged in by an animate being or a natural force." We use Allen's *ACAUSE* predicate to characterize an *action* as a particular kind of event, namely those which have an animate-being as agent.

$$\forall x \text{ action}(x) \leftrightarrow \exists e \text{ event}(x) \wedge \text{eclass}(x) = \text{ACAUSE}(\text{agent}(x), e)$$

Otherwise, if an event is a *change-of-state*, this implies that there is no action which causes the event.

The category *reaction* is described by Rumelhart as "the response of a willful being to a prior event." In our characterization of *reaction*, we take Rule 5 as a guide and define it as an event having two components: an internal response and an external response.

$$\forall x \text{ reaction}(x) \rightarrow \text{event}(x) \wedge \text{class}(x) = \text{and}(\text{internal-response}(x), \text{externalresponse}(x))$$

Internal-response and *external-response* are functions over reactions. The domains of these functions correspond to Rumelhart's syntactic categories *internal response* and *overt response*.

Rumelhart describes the syntactic category *attempt* as "the formulation of a plan and application of that plan for obtaining a desire." It appears in rules 7, 8, and 10. We believe the underlying intuition here is that *attempt* characterizes the action(s) that an agent may take in response to some motivating event. A complication is that *attempt* includes not only the "physical" actions that a character performs, but also the "mental" act of planning those actions.

Clearly, the notion of a plan is central to relating a sequences of actions performed by some character. However, the concept is only vaguely defined by Rumelhart. He is not using the term "plan" in the classical AI sense of a "recipe for action." Rumelhart informally defines *plan* as "the creating of a subgoal which if achieved will accomplish a desired end." When he says "plan" he means more what Pollack [1990] terms a "complex mental attitude".

We move toward a refinement of this interpretation with the following axioms. First, we postulate a two-place predicate *applicable(plan,goal)*, which is true if the *plan*, successfully carried out, will cause the state described by *goal* to hold. Furthermore,

$$\forall p [p : \text{plan}], \forall g [g : \text{goal}] \text{ applicable}(p, g) \rightarrow \text{goal-owner}(g) = \text{agent}(p)$$

where *goal-owner* and *agent* are functions performing the obvious mappings. In the axiom 1, *animate-being*, *plan*, *time-interval*, and *application* are single-position

$$\begin{aligned}
& \forall x [x : \text{attempt}] \text{animate-being}(\text{agent}(x)) \\
& \quad \wedge \text{plan}(\text{attempt-plan}(x)) \wedge \text{time-interval}(\text{time}(x)) \wedge \text{application}(\text{attempt-application}(x)) \\
& \quad \wedge \text{COMMITTED}(\text{agent}(x), \text{attempt-plan}(x), \text{time}(x)) \\
& \quad \wedge \text{applicable}(\text{attempt-plan}(x), \text{attempt-goal}(x))
\end{aligned} \tag{1}$$

predicates which are true when the argument is an element of the indicated type. The function *attempt-plan* maps from an *attempt* element to the (abstract) plan embodied in the attempt, whereas the function *attempt-application* returns the concrete actions (considered as a unit) which are the “carrying out” of the plan. The function *attempt-goal* maps from an *attempt* to the state description toward which the attempt is aimed at causing to hold. *COMMITTED(a,p,t)* is Allen’s predicate signifying that agent *a* intends to carry out the actions composing plan *p* where time *t* is the time of the *intending*, not of the carrying out of the plan.

Of course, these axioms only constrain what must be true in order for some element to be an *attempt*. Further axioms describe how an attempt is related to the event which leads a character to form the *attempt-goal* in the first place. Other axioms define the relationship of an *attempt* to the *application*, which is “carrying out” phase of an attempt.

One of the most important relationships in this class of stories is that of failed attempts to the surrounding events. Of course, in order to relate a failed attempt to another event, we must first be able to recognize when an attempt has, in fact, failed. Axiom 2 defines failure. There are two cases of plan failure to consider which correspond roughly to a base case and a recursive case. In the base case, a plan fails if every action of the plan has been done and the goal still does not hold. In the recursive case, a plan fails if there is a step which is a subgoal and at some time following the time at which the subgoal was to be achieved, the subgoal still has not been met. This portion of the rule is necessary since it is impossible to talk about a step which is a subgoal having been done if the subgoal was not met.

Relations Rumelhart informally describes six semantic relationships on events. They are: *and*, *allow*, *initiate*, *motivate*, *cause*, and *then*. With the exception of *allows*, we have accounted for all of these in our logic. We have developed axioms which characterize *initiate*, *motivate*, and *cause*; *and* and *then* have been dissolved into the logic.

In general, it is difficult to give precise formulations of these relationships without laying groundwork for which we lack the space in this paper, but we will present here some examples of logical axioms we have developed in order to give the reader a sense of the direction in which we are moving. Rumelhart describes the semantic relationship *initiates* as follows: “the re-

lationship between an external event and the willful reaction of an anthropomorphized being to that event.” Using Allen’s primitives for relating time intervals, we construct the predicate *starts-before* as follows:

$$\begin{aligned}
& \forall t_1, t_2 \text{ starts-before}(t_1, t_2) \leftrightarrow \\
& \quad \oplus [\text{in}(t_2, t_1), \text{before}(t_1, t_2), \text{meets}(t_1, t_2), \text{overlaps}(t_1, t_2)]
\end{aligned}$$

We then use this to make explicit at least part of what must be true when *x* *initiates* *y*:

$$\begin{aligned}
& \forall x, y \text{ initiates}(x, y) \rightarrow \\
& \quad \text{starts-before}(\text{time}(x), \text{time}(y)) \wedge \text{reaction}(y)
\end{aligned}$$

Rumelhart uses the predicate *motivate* to describe the relationship between “an internal response and the actions resulting from that response.” We tighten this by means of a predicate *motivates (mental-event, action)* which is true when the mental-event such as an emotion or the adoption of a goal is the motivation for some external action, which may be a high level action composed of or generated by one or more sub-actions. For example, the internal-response component of a reaction to an event *motivates* the agent’s overt-response:

$$\begin{aligned}
& \forall x [x : \text{reaction}] \\
& \quad \text{motivates}(\text{internal-response}(x), \text{overt-response}(x))
\end{aligned}$$

Rumelhart defines the semantic relationship *cause* as “the relationship between two events in which the first is the physical cause of the second.” Like Allen, we take causality to be a primitive relationship and leave it unanalyzed. Rumelhart’s two remaining semantic relationships, *and* and *then* essentially are dissolved into the temporal logic. The former was nothing more than logical conjunction; and the latter was meant to suggest temporal ordering of events, which may be more precisely expressed with Allen’s predicates.

Conclusions

In this paper we provide part of the logical foundations for “cashing out” a story grammar like Rumelhart’s into a generative model of simple narratives. We have provided examples of definitions which specify the “syntactic” categories of Rumelhart’s story grammar and the constraints between constituents (the annotations to the rules).

When this work is completed, we intend to move on to implementing the generative model and using it to create a corpus of stories and their “meanings.” A first step will be to generate a set of event sequences from a story grammar (revised to take into account the newly-formalized constraints) and a body of background knowledge.

$$\begin{aligned} & \forall p [p : \text{plan}], \forall t [t : \text{timeinterval}] \text{failed}(p, t) \leftrightarrow \\ & (\forall a, t' (\text{TODO}(a, t', p) \rightarrow [\text{OCCUR}(a, t') \wedge \text{BEFORE}(t', t)]) \wedge \neg \text{HOLDS}(\text{goal}(p), t)) \\ & \vee \exists sg, t' (\text{TODO}(\text{achieve}(sg), t', p) \wedge \text{BEFORE}(t', t) \wedge \neg \text{HOLDS}(sg, t')) \end{aligned} \quad (2)$$

The most important problem remaining to our formalization project is that of characterizing the relationship of story settings to the corresponding sequences of events. Rumelhart uses the *allows* predicate to express the notion that the states composing the setting set the stage for what follows. That the setting simply not contradict the events which follow is clearly not a strong enough constraint on the setting-body relation. A relevance relationship like this one is difficult to formalize in a framework like the first order logic. We are currently working on characterizing the kinds of facts that are expressed in story settings.

References

- [Allen, 1984] James F. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123-154, 1984.
- [Birnbaum, 1991] Lawrence Birnbaum. Rigor mortis: a response to Nilsson's "logic and artificial intelligence". *Artificial Intelligence*, 47:57-77, 1991.
- [Black and Wilensky, 1979] John B. Black and Robert Wilensky. An evaluation of story grammars. *Cognitive Science*, 3:213-230, 1979.
- [Bower, 1976] Gordon H. Bower. Experiments on story understanding and recall. *Quarterly Journal of Experimental Psychology*, 28:511-534, 1976.
- [Goldman and Lang, 1992] Robert P. Goldman and R. Raymond Lang. Representing intentional actions in allen's temporal logic. forthcoming, 1992.
- [Johnson and Mandler, 1980] Nancy S. Johnson and Jean M. Mandler. A tale of two structures: Underlying and surface forms in stories. *Poetics*, 9:51-86, 1980.
- [Kautz and Allen, 1986] Henry Kautz and James Allen. Generalized plan recognition. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 32-38, 1986.
- [Lebowitz, 1985] Michael Lebowitz. Story telling and generation. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*, pages 100-109, Berkeley, California, 1985.
- [Mandler and Johnson, 1977] Jean M. Mandler and Nancy S. Johnson. Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*, 9:111-151, 1977.
- [McDermott, 1982] Drew V. McDermott. A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101-155, 1982.
- [McDermott, 1987] Drew V. McDermott. A critique of pure reason. *Computational Intelligence*, 3:151-160, 1987.
- [Meehan, 1976] James R. Meehan. *The Metanovel: Writing Stories by Computer*. PhD thesis, Yale University, 1976.
- [Pollack, 1990] Martha E. Pollack. Plans as complex mental attitudes. In Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, editors, *Intentions in Communication*, pages 77-104. MIT Press, 1990.
- [Polti, 1921] Georges Polti. *The Thirty-Six Dramatic Situations*. James Knapp Reeve, Franklin, Ohio, 1921.
- [Propp, 1968] V. Propp. *Morphology of the Folktale*. University of Texas Press, 1968.
- [Rumelhart, 1975] David E. Rumelhart. Notes on a schema for stories. In Daniel G. Bobrow and Allan Collins, editors, *Representation and Understanding: Studies in Cognitive Science*, pages 211-236. Academic Press, Inc., New York, 1975.
- [Schubert and Hwang, 1989] Lenhart K. Schubert and Chung Hee Hwang. An episodic knowledge representation for narrative texts. In *First International Conference on Principles of Knowledge Representation and Reasoning*, pages 444-458. Morgan Kaufmann Publishers, Inc., 1989.
- [Stein and Glenn, 1979] Nancy L. Stein and Christine G. Glenn. An analysis of story comprehension in elementary school children. In Roy O. Freedle, editor, *New Directions in Discourse Processing*, volume 2, pages 53-120. Ablex Publishing Corporation, Norwood, New Jersey, 1979.
- [Thorndyke, 1977] Perry W. Thorndyke. Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology*, 9:77-110, 1977.
- [Turner and Dyer, 1985] Scott R. Turner and Michael G. Dyer. Thematic knowledge, episodic memory and analogy in minstrel, a story invention system. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*, pages 371-375, Berkeley, California, 1985.
- [Wilensky, 1982] Robert Wilensky. Points: A theory of the structure of stories in memory. In W. G. Lehnert and M. H. Ringle, editors, *Strategies for Natural Language Processing*, pages 345-374. Lawrence Erlbaum, Hillsdale, N.J., 1982.