Erik T. Mueller

Computational Models of Narrative

This article provides an overview of work on computational modeling of narrative, with a focus on narrative understanding. It reviews representations for narrative, systems for narrative understanding, and systems for narrative generation. It proposes a major project for narrative understanding and specifies five problems that the project must address: efficiency of reasoning, effective model finding, representation of knowledge for narrative, acquisition of knowledge for narrative, and acquisition of annotated training data.

1 Introduction

Narratives are part of the fabric of our lives. What are they? How do they work? One way to answer these questions is to formulate theories of narrative and discuss them in natural language. But to a computer scientist, this isn't satisfactory. It isn't precise enough. It doesn't lead to a running computer program. In fact, computer scientists tend to believe (and I tend to agree) that we don't truly understand something until we build a computer program that does it. Writing a computer program forces us to explore all the gaps in our understanding. If our understanding is gibberish, then the program will crash.

Goals of CMN. The goal of computational modeling of narrative (CMN) is to build computer programs that understand and generate narratives.

This article focuses on models for the understanding of stories; for an overview of story generator algorithms, see the eponymous article by Gervás (2012). However, representations such as the ones described in Section 2 are common to narrative understanding and narrative generation. Generation may require further representations for dealing with writing style and dramatic goals such as suspense.

In the past, there has been more work on computational modeling of narrative understanding than on computational modeling of narrative generation, although this is starting to change, especially with the work on interactive narrative (cf. the notes on *interactive narrative* below, section 4).

In one sense, it is easier to write a narrative generation program than a narrative understanding program, because a human interpreter will read more into a generated narrative than is actually there. Current narrative generation programs can generate many paragraphs of plausible text, whereas current narrative understanding programs have difficulty understanding narratives as short as three sentences. In another sense, it is harder to write a narrative generation program because generating good narratives requires creativity (for a classical discussion and implementation, cf. Turner, 1994). Narrative generation requires deciding what the narrative is to be about, whereas narrative understanding only requires that the listener or reader understand what has already been written by someone else.

This separation of narrative understanding and generation is paralleled in computational linguistics. Natural language parsing and natural language generation are very different fields (cf., e.g., Dale & Reiter, 2000).

Given a narrative, a program for narrative understanding should be able to explain who the characters are, what they want, where they are, what they do, why they do what they do, what obstacles they face, what happens, and what the point of the narrative is. The program should also be able to construct narratives that resemble those a person might construct.

CMN began in the early 1970s in artificial intelligence laboratories and computer science departments at the Massachusetts Institute of Technology, Yale University, and elsewhere. From the start, researchers in this area have had a great interest in how humans represent and process narratives, and there has been significant cross-pollination between AI researchers and psychologists who study human narrative processing. As AI became more empirical and focused on experimental evaluation in the 1990s, CMN also became more empirical. In recent years, there has been increased involvement from literary theorists, narratologists, and neuroscientists.

Relationship with Computational Linguistics. How does computational modeling of narrative differ from computational linguistics and natural language processing? Of course, narrative is a kind of discourse, which is one of the domains of computational linguistics and natural language processing. But CMN differs in two main ways. First, CMN focuses on processing and representation of areas that sometimes seem far removed from computational linguistics like emotions, personality traits, counterplanning, plot structures, and story themes. Modeling narrative requires addressing these areas. Second, CMN focuses on building complete, integrated computer systems that perform high-level cognitive tasks like creating a story from scratch or asking probing questions to resolve gaps in understanding. Building such systems is difficult, but working on them generates new problems (like how to process stereotypical situations) that drive the larger field of natural language processing forward.

Related Publications. Previous summaries and discussions of computational modeling of narrative are provided by Mueller (2000b), Mueller (2002), and Richards, Finlayson, and Winston (2009). Book-length treatments of CMN are provided by Ram and Moorman (1999) and Mani (2013). Books on narrative and narrative understanding include those by Prince (1982), van Dijk and Kintsch (1983), Ryan (1991), Emmott (1997), Kintsch (1998), S. R. Goldman, Graesser, and van den Broek (1999), and Rimmon-Kenan (2002). Lehnert (1994) provides a memoir of her experiences in CMN. Riloff (1999) discusses the relation between information extraction and narrative understanding. More detail on knowledge representation is provided by Davis (1990), Lenat and Guha (1990), Brachman and Levesque (2004), and van Harmelen, Lifschitz, and Porter (2008).

This article is focused on computer science and artificial intelligence modeling of narrative. Narrative has also been extensively studied in the field of psychology. A good summary is given by Zwaan and Radvansky (1998). Seminal works in this area are: H. H. Clark (1977), Just and Carpenter (1980), Kintsch and van Dijk (1978) and Kintsch (1988). McKoon and Ratcliff (1986, 1992) and Graesser, Singer, and Trabasso (1994) present important controversies and experimental psychological evidence. Some narrative processing systems have been based directly on psychological models of narrative; interesting examples are those by Mross and Roberts (1992) and by Langston, Trabasso, and Magliano (1999).

Workshop series on CMN. Two workshop series on CMN are regularly held: the workshops on *Computational Models of Narrative,*[1] behind which Mark Finlayson is the driving force, and the workshops on *Intelligent Narrative Technologies*[2] run by Mark Riedl and Brian Magerko.

The regularly held *SemEval* evaluations, which started as *Senseval* in Sussex, UK, in 1998, in computational linguistics have several tasks related to narrative understanding, including word sense disambiguation, semantic role labeling, recognizing textual entailment, temporal annotation, and sentiment analysis (cf. the latest volume, Manandhar & Yuret, 2013).

2 Representations for Narratives

People read and answer questions about a narrative by building, examining, and manipulating mental representations (Graesser et al., 1994; Kintsch, 1988; Zwaan & Radvansky, 1998). We can implement these representations as data structures and build programs that process them in a way similar to how humans do it.

---

1 The first workshop was the *MIT Workshop on Computational Models of Narrative*; it took place in 2009 in Beverley, Massachusetts, USA, and is documented by Finlayson, Richards, and Winston (2010); CMN 2010 (Finlayson, 2010) was the *2010 AAAI Fall Symposium on Computational Models of Narrative* in Arlington, Virginia, USA; CMN 2012 (Finlayson, 2012) took place as a workshop of the *Language Resources and Evaluation Conference* 2012, İstanbul, Turkey; CMN 2013 (Finlayson, Fisseni, Löwe, & Meister, 2013) was colocated with the *36th Annual Meeting of the Cognitive Science Society* (CogSci 2013) and took place in Hamburg, Germany; CMN 2014 will be colocated with CogSci 2014 in Québec City, Canada.

2 INT1 (Magerko & Riedl, 2007) took place at the *AAAI 2007 Fall Symposium*, in Arlington, Virginia, USA; INT2 (Louchart, Mehta, & Roberts, 2009) at the *AAAI 2009 Spring Symposium* in Stanford, California, USA; INT3 (Jhala, Riedl, & Roberts, 2010) at the *Foundations of Digital Games* conference 2010 in Monterey, California, USA; INT4 (Tomai, Elson, & Rowe, 2011) at the *Seventh Annual AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (AIIDE) 2011 in Palo Alto, California, USA; INT5 (Ware, Zhu, & Hodhod, 2012) at AIIDE 2012, Stanford, California, USA; INT6 (Cavazza, Si, & Zook, 2013) at AIIDE 2013, Boston, Massachusetts, USA; INT7 will be colocated with the *Electronic Literature Organization Conference 2014* in Milwaukee, Wisconsin, USA.

Now, it may seem futile to try to pin down mental representations for narratives, when narratives may discuss anything under the sun and may even be about a topic that the reader knows nothing about. Also, the narrative may use words and phrases that are unknown to the reader. By reading the narrative, the reader is able to learn about the topic and learn the meanings of the new words and phrases.

But certain realms repeatedly occur in narratives. If we can build computational models of these realms, we can go a long way toward getting computers to process narratives.

We start with a realm that is often addressed in computational linguistics, namely the realm of states and events involving people and things. We then proceed toward realms that are progressively less in the purview of computational linguistics.

2.1 People, Things, States, and Events

Consider this narrative fragment: *Sophia was amused by the cat. She laughed.* Here we have a person Sophia, a state, Sophia being amused by the cat, and an event, Sophia laughing.

We represent states and events using predicate-argument structure (Kingsbury & Palmer, 2002) or first-order logic (Ebbinghaus, Flum, & Thomas, 1994). The representation of the above sentences is something like **{AmusedBy(Sophia1, Cat1), Laugh(Sophia1)}**. **AmusedBy(Sophia1, Cat1)** represents a state, and **Laugh(Sophia1)** represents an event. The word "state" here refers to the state of a person or thing in the world. In some formalisms like state transition systems (Ghallab, Nau, & Traverso, 2004), "state" refers to the entire state of the world.

There are two contrasting schools of thought on predicates like **AmusedBy** and **Laugh**. One school attempts to represent all concepts using a small set of *semantic primitives* (Dorr, 1993; Jackendoff, 1972, 1990; Schank, 1972). The other school uses as many predicates as necessary, perhaps as many as one predicate for every concept (Viegas, 1999).

The arguments for a particular predicate are sometimes given names, called *roles*, like **agent**, **patient**, **object**, **source**, **destination**, **beneficiary**, **recipient**, and **instrument** (Baker, Fillmore, & Lowe, 1998; Fillmore, 1968; Palmer, Gildea, & Kingsbury, 2005). Thus we may represent *Sophia went into the living room* as **Go(agent: Sophia, destination: LivingRoom1)**.

A number of *semantic parsers* (Alshawi, 1992; Blackburn & Bos, 2005; McCord, Murdock, & Boguraev, 2012) and *semantic role labelers* (Gildea & Jurafsky, 2002; Punyakanok, Roth, & Yih, 2008) have been developed for converting text into predicate-argument structure. Once we have the predicate-argument structure, we use it to answer questions. Given the question *Who was amused by the cat?*, we use a semantic parser to convert the question into **AmusedBy(?answer, Cat1)**. We then match this against the representation of the narrative, which yields **Sophia1** as the value of **?answer**. We then convert **Sophia1** back into text, which yields *Sophia*, the answer to the question.

2.2 Preconditions and Effects of Events

Consider *Sophia took the mitten from the cat.* From this we infer that (1) before Sophia took the mitten from the cat, the cat had the mitten, and (2) after Sophia took the mitten, Sophia had the mitten. For a computer program to make these inferences, it must have knowledge about the *preconditions* and *effects* of events (Fikes & Nilsson, 1971; Ghallab et al., 2004; Newell & Simon, 1961).

The preconditions of an event are requirements that must be satisfied for an event to occur. The effects of an event are how the world is changed by the event. For the event **Take(?agent, ?object, ?source)**, the precondition is **Have(?source, ?object)**. For an agent to take an object from a source, the source must have that object in the first place. The effects of this **Take** event are **not Have(?source, ?object)** and **Have(?agent, ?object)**. After the agent takes the object from the source, the source no longer has the object, and the agent has the object.

Given the above sentence, the semantic parser produces **Take(Sophia1, Mitten1, Cat1)**. Using the above preconditions and effects, we infer that **Have(Cat1, Mitten1)** holds before the event and that **not Have(Cat1, Mitten1)** and **Have(Sophia1, Mitten1)** hold after the event. Mueller (2006) provides a detailed treatment of reasoning about events using the event calculus (Miller & Shanahan, 2002).

2.3 Time

Narratives involve states and events that unfold over time: *Sophia opened the front door, and then she went into the living room.* A narrative understanding program must represent that the first event occurred before the second event.

One simple representation for time is to use an ordered list of states and events: **[Open(Sophia1, FrontDoor1), Go(Sophia1, LivingRoom1)]**. But this does not allow us to specify overlapping states and events. To do this, *temporal relations* are usually introduced. Allen (1983) defines the following temporal relations: **BEFORE**, **EQUAL**, **MEETS**, **OVERLAPS**, **DUR- ING**, **STARTS**, and **FINISHES**. The TimeML standard (TimeML Working Group, 2005) defines the following temporal relations: **BEFORE**, **AFTER**, **ON_OR_BEFORE**, **ON_OR_AFTER**, **LESS_THAN**, **MORE_THAN**, **EQUAL_OR_LESS**, **EQUAL_OR_MORE**, **START**, **MID**, **END**, and **APPROX**.

Temporal relations can be used in several ways. One way is to assert temporal relations between events (Hobbs, Stickel, Appelt, & Martin, 1993). In this scheme we represent the above narrative as **BEFORE(Open(Sophia1, FrontDoor1), Go(Sophia1, LivingRoom1))** or as **{E1=Open(Sophia1, FrontDoor1), E2=Go(Sophia1, LivingRoom1), BEFORE(E1, E2)}**.

Another way is to represent time using the real number line. We assert that events happen and states hold over time intervals, and then we assert temporal relations between time intervals (Allen, 1984). In this scheme we represent the above narrative as **{Happens(Open(Sophia1, FrontDoor1), T1), Happens(Go(Sophia1, LivingRoom1), T2), BEFORE(T1, T2)}**.

Other possibilities include considering time to be a sequence of states in a state transition system (Ghallab et al., 2004) and considering time to be a branching real number line or a tree, which allows reasoning about alternative hypothetical sequences of events (McCarthy & Hayes, 1969; Reiter, 2001). Complete reviews of temporal models are provided by Gabbay and Ohlbach (1994) and Fisher, Gabbay, and Vila (2005). ter Meulen (1995) discusses representations for time in natural language.

The challenge is to construct a representation of the time course of the states and events given the input narrative. In a narrative, the default is for events and states to be presented in sequence. When events and states are presented out of sequence, this is sometimes indicated using verb tense or aspect and temporal phrases like *before*, *after*, *during*, *while*, and *the previous day*. But this is not always the case. Consider the following: *Sophia got out of the pool. The water was too cold.* Here the water was cold before Sophia got out of the pool. Knowledge about people's reactions to the temperature of water is needed to make this inference.

2.4 Space

Narratives also deal with space. Consider the following: *Sophia went into the kitchen. She took the scallions from the fridge and set them on the cutting board.* A narrative understanding program must track the locations of people and things over time (Duchan, Bruder, & Hewitt, 1995).

Dyer (1983) presents the *scenario participant map* for tracking characters in space and time. It consists of a graph of settings (like a hotel room) connected by transitions (like walking through a hallway to get from one room to another). Davis (1990, 1995) discusses a number of representations of space. Randell, Cui, and Cohn (1992) propose a logic for spatial reasoning. Morrow (1994) discusses creating spatial models from text. Mueller (1998) discusses the use of two-dimensional grids for representing typical locations like a grocery store, theatre, and hotel room, and for tracking characters in space and time. Kuipers (2000) describes the spatial semantic hierarchy, a collection of interacting representations of large-scale space. Gerard and Sansonnet (2000) present a spatiotemporal representation for narratives. Slobin (2003) treats motion events.

2.5 Stereotypical Situations and Scripts

Consider the following: *Sophia sat down in the front row. The teacher entered.* We recognize that Sophia is probably a student who is attending a class. A narrative understanding program must recognize and classify *stereotypical situations*. The program must also fill in missing events based on the stereotypical situation. For example, it should infer that, assuming everything went as planned, the teacher taught the class and, once the class was over, Sophia left the classroom.

Schank and Abelson (1977) pointed out the importance of these situations, which they call *scripts*, in narrative understanding and, by my count, identified 51 scripts that commonly occur in narratives. A script consists of (1) a set of *roles* and (2) a graph of *events* that reference these roles. For example, a simplified version of the **RESTAURANT** script would contain the roles **customer, table, menu, waiter, food,** and **check,** and the following events: (1) **customer** enters, (2) **customer** sits at **table**, (3) **customer** reads **menu**, (4) **customer** orders, (5) **waiter** brings **food**, (6) **customer** eats **food**, (7) **customer** pays **check**, and (8) **customer** leaves. Schank and Abelson (1977) also discuss two kinds of script deviations: *interferences*, in which a script does not proceed normally because an action precondition is not satisfied or an action is performed incorrectly, and *distractions*, in which unexpected events initiate new goals that cause the script to be abandoned.

Gordon (1999, 2000) developed a database of 768 activities. Each activity contains the following slots: (1) **Events**, (2) **Places**, (3) **People**, (4) **Things**, and (5) **Misc** (entry conditions and results). Slot values are taken from the United States Library of Congress Thesaurus for Graphic Materials (Library of Congress, 1995). Mueller (2000a) built predicate-argument structure representations of 100 scripts for the ThoughtTreasure narrative understanding system. Mueller (2004, 2007a) built event calculus representations of 15 scripts: the restaurant script, four terrorism scripts, and 10 scripts frequent in American literature texts.

Script classification can be treated as a text categorization problem and solved using statistical natural language processing techniques (Manning & Schütze, 1999). The task is to assign a segment of a narrative to one of many scripts. Chambers and Jurafsky (2008, 2009) developed an unsupervised learning algorithm for learning new scripts, including roles.

2.6 Goals and Plans

Understanding the behavior of characters in a narrative requires recognizing the *goals* of the characters and their *plans* for achieving these goals. Consider the following: *Sophia loved the pearl necklace. She got out her credit card.* We infer that Sophia got out her credit card because this was part of her plan for achieving her goal to own the pearl necklace.

Schank and Abelson (1977) provide a taxonomy of human goals that commonly appear in narratives. *Satisfaction goals* are goals to satisfy recurring needs: **S-HUNGER** (satisfy hunger), **S-SEX** (satisfy a need for sex), and **S-SLEEP** (satisfy a need for sleep). *Enjoyment goals* are goals pursued for enjoyment: **E-TRAVEL** (enjoy travel), **E-ENTERTAINMENT** (enjoy entertainment), and **E-COMPETITION** (enjoy competition). *Achievement goals* are goals to achieve or acquire something: **A-GOOD-JOB** (achieve a good job), **A-POSSESSIONS** (acquire possessions), and **A-SOCIAL-RELATIONSHIPS** (acquire social relationships). *Preservation goals* are goals to preserve or maintain something: **P-HEALTH** (preserve health), **P-JOB** (preserve job), **P-POSSESSIONS** (preserve possessions), and **P-SOCIAL-RELATIONSHIPS** (preserve social relationships). They also describe goals that involve imminent threats (**C-HEALTH, C-FIRE,** and **C-STORM**) and goals that are instrumental to other goals (**I-PREP**, **D-KNOW**, **D-PROX**, and **D-CONT**).

Schank and Abelson also presented plans for achieving goals. For example, **USE(FOOD)** is a plan for the **S-HUNGER** goal. This plan consists of the following steps: (1) knowing (**D-KNOW**) the location of some food, (2) being near (**D-PROX**) the food, (3) having control (**D-CONT**) over the food, and (4) eating the food.

A narrative understanding system must track the evolution of the goals and plans of the characters over time, including activation of goals, activation of plans for achieving goals, execution of actions on behalf of plans, and success or failure of plans and goals. Narrative understanding also requires tracking the relationships among multiple goals (Wilensky, 1983). Goals may be compatible or in conflict, both within a character and among characters. Carbonell (1980) shows how human personality traits can be modeled in terms of goal priorities.

2.7 Emotions and Sentiments

Understanding narratives requires understanding *emotions* as well as *sentiments* or the way a character feels about a person or thing. Consider the following: *Sophia opened the present from Emma. It was a beautiful scarf! She was grateful.* We infer that Sophia liked the scarf and that she was grateful to Emma for giving her the scarf.

Dyer (1983) pointed out the importance of emotions in narrative understanding and proposed a representation of emotions useful for narrative understanding. The representation consists of a **STATE** (positive or negative), **CHAR** (character experiencing the emotion), **G-SITU** (goal situation leading to the emotion), **TOWARD** (person toward whom the emotion is directed), **SCALE** (intensity level), and **E-MODE** (whether the goal situation is expected or unexpected).

This scheme enables representation of a number of emotion words that occur in narratives. *Happy*, *joyous*, and *glad* are positive emotions associated with goal success. *Unhappy*, *upset*, and *sad* are negative emotions associated with goal failure. *Grateful* and *thankful* are positive emotions toward a person who caused a goal success for **CHAR**. *Annoyed*, *angry*, and *furious* are negative emotions toward a person who caused a goal failure for **CHAR**. *Hopeful* is a positive emotion associated with an active goal that is expected to succeed. *Fearful* and *worried* are negative emotions associated with an active goal that is expected to fail. *Surprised* is a positive emotion associated with an unexpected goal success. *Shocked* is a negative emotion associated with an unexpected goal failure. *Relieved* and *allayed* are positive emotions associated with a preservation goal that succeeds. *Disappointed* is a negative emotion associated with an unexpected goal failure. *Proud* and *smug* are positive emotions associated with causing a goal success for another person. *Guilty*, *ashamed*, *embarrassed*, and *regretful* are negative emotions associated with causing a goal failure for another person.

Ortony, Clore, and Collins (1988) present a representation of emotions that was used in the AbMaL narrative understanding system (O'Rorke & Ortony, 1994). Heider (1958) discusses sentiments in detail, and describes tendencies in sentiments, such as the fact that people tend to like people who have similar sentiments. Pang and Lee (2008) review techniques for sentiment analysis.

2.8 Themes and Plots

So far we have discussed a number of realms that a narrative understanding program must track in a narrative, including states, events, space, time, goals, and emotions. In addition, there is usually some reason why the writer or speaker is telling the narrative, or there is some overarching theme, point, or imperative that the writer or speaker is trying to convey to the reader or listener. Consider the following: *Sophia was always bragging about herself and never paying attention to anyone around her. One day her car broke and she needed a ride to work, but nobody would help her.* A person reading this easily extracts the moral, *Be kind to others so that they will be kind to you.* Dyer (1983) developed a representation of narrative adages and morals called *thematic abstraction units* (TAUs). A TAU consists of (1) a plan, (2) the intended effect of the plan, (3) why the plan failed, and (4) how to avoid such a failure in the future. For example, the adage *Don't close the barn door after the horse has escaped* may be represented as the following TAU, called **TAU-POST-HOC**:

(1) a plan P (*close the barn door*) to achieve preservation goal G (*keep the horse in the barn*), (2) P is intended to satisfy enablement condition C (*barn door closed*) for G, (3) the planner failed to execute P while G was active and C was unsatisfied, and the planner executed P after G failed, and (4) the planner should execute P while G is still active and C is unsatisfied, and the planner should not execute P after G fails.

Turner (1994) developed *planning advice themes* (PATs), which extend TAUs with a more detailed representation of planning advice. A PAT consists of the following: (1) **Decision Point** (whether the PAT is invoked during plan selection or goal activation), (2) **Value** (whether the PAT provides positive or negative advice), (3) **Decision** (the planning choice), (4) **Consequence** (the consequence of the **Decision**), (5) **Connection** (the causal link between the **Decision** and the **Consequence**), (6) **Object** (the specific object within the **Decision** to which the advice applies), (7) **Planner** (the specific character to which the advice applies), (8) **Current Goal** (the current goal of the **Planner**), (9) **Current Plan** (the current plan of the **Planner**), (10) **Active Goals** (the active goals of the **Planner**), and (11) **World Facts** (facts that must be true for the advice to apply).

Lehnert (1982) developed *plot units* for summarizing narrative plots. A plot unit representation of a narrative is a graph containing three types of nodes: **M** (motivational states like active goals), + (positive states like goal successes), and - (negative states like goal failures). The graph contains a chronological list of states for each character. The states are related via cross-character links and intra-character links: **a** (actualization), **t** (termination), **e** (equivalence), and **m** (motivation). Because plot unit graphs represent only abstract states and not the particular type of goals, plot units provide a more abstract representation of a narrative than a goal-plan analysis. For example, the

RETALIATION plot unit describes any narrative in which character A causes a negative state for character B, which causes character B to cause a negative state for character A. Goyal, Riloff, and Daumé (2010) present techniques for automatically producing plot unit representations from narrative text. Nackoul (2010) presents English patterns for locating instances of plot units in text.

Other representations for narrative themes include *thematic organization packets* (TOPs) (Schank, 1982), *story points* (Wilensky, 1982), and *story intention graphs* (SIGs) (Elson, 2012).

3 Narrative Understanding Systems and Algorithms

Methods for building narrative understanding systems have been continually evolving since the first systems were developed in the early 1970s. I discuss systems and algorithms for narrative understanding in roughly chronological order.

3.1 Systems based on Rules, Demons, and Agents

An early narrative understanding system was built by Charniak (1972) at MIT. The system is based on condition-action pairs or IF-THEN rules called *demons* (Charniak, Riesbeck, & McDermott, 1980; Minsky, 1961; Selfridge, 1959). Sample demons are: (1) **IF** it is raining and person P is outside, **THEN** assert that P is wet, and (2) **IF** person P shakes piggy bank B, and money M comes out of B, **THEN** assert that M comes out of B because P shakes B. The system was able to handle two narrative fragments that had been manually converted into the system's internal representation. Charniak (1977a, 1977b) later developed a narrative understanding system called Ms. Malaprop based on frames (Minsky, 1974). Rosenberg (1977) developed a system for understanding news articles using frames.

Other early natural language understanding programs were developed by Winograd (1972) and Schank, Goldman, Rieger, and Riesbeck (1973, 1975).

Starting in the mid-1970s, Roger Schank and his students at Yale developed a number of narrative understanding systems. They are knowledge-intensive systems, making use of scripts, plans, goals, and other knowledge structures (Schank & Abelson, 1977). The knowledge is applied using various algorithms as well as demons, sometimes also called *predictions*, *expectations*, *experts*, and *requests*. Schank and Riesbeck (1981) present Lisp code for micro versions of several of the systems developed during this time.

Cullingford (1978) developed the Script Applier Mechanism (SAM), which uses scripts to understand news stories about five scripts: motor vehicle accidents, plane crashes, train wrecks, oil spills, and state visits. SAM uses several types of patterns, called *script headers*, to activate scripts. A *locale header* activates a script when a character moves to the setting of the script. A *precondition header* activates a script when the narrative mentions an entry condition of the script. An *instrumental header* activates two scripts when a character uses one of the scripts as an instrument for the other script. A *direct header* activates a script when the narrative simply states that the script occurred. Lehnert (1978) developed a model of question answering and associated program QUALM that was used by SAM. DeJong (1979) developed a system called FRUMP that was able to handle 50 scripts.

Wilensky (1978) developed the Plan Applier Mechanism (PAM), which uses plans and goals for narrative understanding and also applies knowledge of goal interactions like conflict, concord, and subsumption. PAM's knowledge was represented as 180 requests (or demons) such as **IF** person P has the goal to possess object O, **THEN** P probably wants to use O for something. PAM was able to process 16 classes of short narratives.

Carbonell (1979) developed the POLITICS system, which uses goal trees and counterplanning strategies to model interpretation of international events from the point of view of different political ideologies.

Dyer (1983) developed the BORIS system, which incorporates many of the knowledge structures used in previous systems, including scripts, plans, and goals, and extends them with additional knowledge structures for memory organization packets (or MOPs, generalizations of scripts), spatiotemporal organization, emotions, interpersonal relations, and narrative morals. BORIS is an integrated system in which knowledge structures of all types are coded as demons, and demons are able to use information from other demons as a narrative is being processed. BORIS was able to perform in-depth understanding of two narratives about divorce.

Alvarado (1990) built OpEd, an editorial understanding system. Reeves (1991) built THUNDER, which handles narratives involving irony. August (1991) developed ARIEL, a system for understanding analogies in arguments.

Norvig (1987) built FAUSTUS, which uses marker passing to unify multiple inference methods. Palmer, Passonneau, Weir, and Finin (1993) built KERNEL, which addresses integration of multiple knowledge sources during text understanding. Mahesh (1995) built COMPERE, which addresses syntax-semantics interaction in sentence understanding.

Minsky (1986) proposed that narrative understanding is carried out by multiple agents operating in different realms. For example, *Mary gave John the kite* is understood simultaneously by agents concerned with physical, possessional, and psychological realms. Mueller (1998) built ThoughtTreasure, which uses agents to build models of a narrative. McCarthy et al. (2002) discuss integration of multiple methods for narrative understanding.

3.2 Case-Based Systems

As an approach to scaling up narrative understanding systems and allowing them to learn, students of Schank developed a number of memory-based narrative understanding systems (Schank, 1982). Kolodner (1984) developed CYRUS, a memory system for FRUMP, and Lebowitz (1980) developed the IPP system. These systems make generalizations based on narratives previously processed and use these generalizations to help understand new narratives. This idea evolved into *case-based reasoning* (Kolodner, 1993;

Leake, 1996; Riesbeck & Schank, 1989). Based on previous work on ROBIN (Lange & Dyer, 1989), Lange and Wharton (1992) built REMIND, which generates interpretations of an input text passage by retrieving similar passages and episodes from memory.

3.3 Explanation-Based Systems

Schank (1986) proposed to make computers more creative by having them ask questions about unexpected or odd situations to produce explanations. Several systems were implemented based on this idea (Schank, Kass, & Riesbeck, 1994).

Ram (1989) built AQUA, a narrative understanding system that identifies gaps in its understanding, uses these gaps to generate questions, answers these questions to produce explanations, and stores these explanations in memory to be used in processing future narratives. AQUA handles several variations of ten types of news stories about terrorism.

Schank's students also developed the explanation-based SWALE system (Schank et al., 1994). Owens (1990) built Retriever and Anon, which address indexing and retrieval of abstract planning knowledge. Leake (1992) built Accepter, which addresses the evaluation of explanations. Kass (1990) built Tweaker and ABE, which address the adaptation of old explanations to new situations. Moorman (1997) built ISAAC, a model of creative reading.

3.4 Systems based on Neural Computation

After the PDP Research Group published an influential two-volume set on connectionism and artificial neural networks (ANNs) (McClelland, Rumelhart, & PDP Research Group, 1986; Rumelhart, McClelland, & PDP Research Group, 1986), CMN researchers began to investigate how to build ANN-based narrative understanding systems. ANNs are networks of nodes and links inspired by brain physiology. ANN representations may be *local* or *distributed*. In a local representation, single nodes represent single concepts. In a distributed representation, a concept is spread across multiple nodes.

Dolan (1989) built CRAM, which takes fable-like narratives as input and produces thematic summaries as output. CRAM combines traditional components for natural language analysis and generation with distributed ANN components for memory storage and retrieval. CRAM handles four narratives: Secretary Search, Professor and Proposal, The Fox and the Crow, and The Bear and the Raccoon.

Miikkulainen (1993) built DISCERN, which reads and answers questions about script-based narratives. DISCERN is built entirely using distributed ANN components. DISCERN handles three scripts (restaurant, shopping, and travel), three script variations or tracks per script, and five roles per script. DISCERN's episodic memory is first trained on a set of artificially generated narratives. The trained episodic memory is then used to fill in missing information and answer questions about new narratives.

St. John (1992) developed the connectionist Story Gestalt model, which uses constraint satisfaction to generate bridging and predictive inferences in narrative understanding. Input layers of the model use local representations, and internal layers use distributed representations. The model is trained on automatically generated texts involving six scripts (beach, airport, restaurant, bar, race, and frisbee). Golden and Rumelhart (1993) developed a local network model of narrative understanding in which a narrative is represented as a trajectory through a situation state space. A situation is specified by a set of propositions each of which has the value **present** or **absent**. Frank, Koppen, Noordman, and Vonk (2003) extended this model with distributed representations of propositions.

Based on previous work on metaphor (Lakoff & Johnson, 1980) and neural schemas (Arbib, Conklin, & Hill, 1987), Narayanan (1997) built KARMA, which uses a representation called *x-schemas*, similar to Petri nets (Reisig, 1985), to represent and reason about motion words like *walk*, *push*, *slide*, and *stumble* in narratives about international economics. KARMA and the neural theory of language on which it is based are also discussed by Feldman (2006) and Lakoff and Narayanan (2010).

3.5 Systems Based on Plan Recognition

Kautz (1991) formalizes plan recognition and presents algorithms for generating explanation graphs given observations of events and a plan library that specifies how events are decomposed into other events. Consider a plan library with the following information: (1) making spaghetti marinara consists of making spaghetti and making marinara sauce, and (2) making eggplant marinara consists of making eggplant and making marinara sauce. Given this plan library and the observation *Ava made marinara sauce*, Kautz's plan recognition algorithms produce a graph in which making spaghetti marinara and making eggplant marinara are two possible explanations of Ava's action. A more detailed plan library would contain multiple levels of event decomposition, specifying, for example, that preparing a meal consists of making spaghetti marinara, eggplant marinara, or other dishes.

Charniak (1983, 1986) built the Wimp system, which uses marker passing for plan recognition. Probabilistic accounts of marker passing have been developed by R. P. Goldman (1990), Carroll and Charniak (1991), and Wu (1992).

3.6 Systems Based on Logic

McCarthy (1959, 1990) proposed to use logic to give computers common sense and to use logic for narrative understanding. Schubert and Hwang (1989, 2000) developed EPILOG, an implementation of episodic logic useful for inferencing in narrative understanding. Dahlgren, McDowell, and Stabler (1989) developed KT, a logic-based inference system for text understanding. Hobbs et al. (1993) developed TACITUS, a system that implements all levels of natural language processing, from the lexicon to script processing, using logic and a general explanation (abduction) engine. Mulkar-Mehta (2000) developed Mini-TACITUS, a Java implementation of TACITUS. Shapiro and Rapaport (1995) applied SNePS, a knowledge representation system based on logic and frames, to narrative understanding. Zarri (1996) developed NKRL, a language for representing the content of narratives. Nossum (2003) developed a logical approach to context in narratives. Löwe, Pacuit, and Saraf (2009) formalized the preferences and beliefs of characters in episodes of a television crime series.

Psychologists have long argued that narrative understanding involves creating mental models of the states and events depicted in the narrative (Bower, 1989; Johnson-Laird, 1983). Mueller (2003) demonstrated how off-the-shelf model finding programs can be used to build narrative models given (1) a declarative representation of the narrative and (2) a declarative representation of the knowledge necessary to understand the narrative. Mueller (2004, 2007a, 2007b) used the model finding approach to build systems to understand narratives involving scripts, plans, and goals. Michael (2010, 2012, 2013) proposed model theoretic definitions of narratives and narrative understanding.

4 Narrative Generation Systems

Compared to narrative understanding, there has been less work on narrative generation. An early narrative generation system was developed by Klein et al. (1973). It uses a simulation language to describe the behavior of characters in different situations. The system generates murder mystery narratives.

Meehan (1976) developed the TALE-SPIN system, which generates English narratives by simulating the goal-based behavior of characters. The narratives result from the input parameters, which include: (1) characters (like bear, bee, fox, crow, and canary), (2) the problem of the main character (like hungry, thirsty, and tired), (3) objects (like berries, flower, and water), (4) the personality traits of the characters (like honesty, kindness, and deceitfulness), (5) the relationships of the characters (like affection, trust, and domination), (6) the beliefs about the traits of other characters, and (7) the initial knowledge states of the characters. Schank and Riesbeck (1981) present Lisp code for a micro version of TALE-SPIN, which was later translated into Common Lisp by Sack (1992). P. Clark (1999) wrote a Prolog version of TALE-SPIN. Cox (1996) evaluated his learning system on narratives automatically generated by TALE-SPIN.

Dehn (1981, 1989) developed the Author narrative generation system, which models the author's process. It incorporates author goals as well as an episodic memory of incidents, characters, and previous narratives to provide material for creating new narratives and distractions to influence the creative process.

Lebowitz (1984, 1985) created UNIVERSE, which generates soap opera plots based on representations of the goals, personality traits, interpersonal relationships, and histories of characters. UNIVERSE contains a library of plans, called plot fragments, for achieving goals. A plot outline is generated by repeatedly selecting a plot fragment for an active goal and executing the plan, which may cause new goals to be activated.

Turner (1994) developed MINSTREL, which is able to generate ten narratives about King Arthur and the Knights of the Round Table. MINSTREL incorporates author goals including thematic goals, dramatic writing goals, consistency goals, and presentation goals. Themes are represented by planning advice themes (PATs). MINSTREL uses four techniques to achieve its dramatic writing goals: suspense, tragedy, characterization, and foreshadowing.

Bringsjord and Ferrucci (2000) developed BRUTUS, an architecture for narrative generation that incorporates domain knowledge, literary knowledge, goals, plans, actions, production rules, and logic. BRUTUS generates narratives about betrayal.

Callaway (2000) developed StoryBook and the Author architecture for generating high-quality narrative prose given a specification from a narrative planner.

Pérez y Pérez and Sharples (2001) developed MEXICA, which generates narratives about the early inhabitants of the Valley of Mexico. Based on sets of narrative actions and previous narratives defined by the user, MEXICA engages in a cycle of engagement and reflection to generate new narratives.

Riedl and Young (2010) developed Fabulist, a narrative generation system that (1) generates a sequence of character actions, (2) generates a discourse plan for expressing the narrative, and (3) realizes the plan.

In 1928, Propp published *Morphology of the Folktale* (Propp, 1968), which presented a framework for characterizing and generating narratives, similar to story grammars (Prince, 1973; Rumelhart, 1975). Fisseni, Kurji, and Löwe (2014) investigated the inter-annotator agreement of annotators using Propp's framework, and Gervás (2013) built a system to generate Russian folk tales using the framework.

Other narrative generation programs include ROALD (Yazdani, 1983), Racter (1984), ALIBI (Kuflik, Nissan, & Puni, 1991), TAILOR (T. C. Smith & Witten, 1991), Dramatica (Phillips & Huntley, 1993), Joseph (Lang, 1997), MAKEBELIEVE (Liu & Singh, 2002), the case-based reasoning system of Gervás, Díaz-Agudo, Peinado, and Hervás (2005), and STORY (Fayzullin, Subrahmanian, Albanese, Cesarano, & Picariello, 2007).

An *interactive narrative* allows the viewer or reader to participate in a narrative as it unfolds (Laurel, 1986). Mueller (1990) developed DAYDREAMER, a computer model of daydreaming that produces narrations of its daydreams in response to interactions with the user. Bates, Loyall, and Reilly (1994) developed the Oz architecture for interactive narrative with goal-based, emotional characters (Loyall, 1997; Reilly, 1996; S. Smith & Bates, 1989; Weyhrauch, 1997). Brooks (1999) developed Agent Stories to support the development of interactive narratives. Mateas and Stern (2003) developed Façade, an interactive narrative in which the user has been invited over to a friend's house for cocktails. Szilas and Rety (2004) developed IDtension, an interactive drama system that incorporates modeling of the audience. Montfort (2003) discusses text-based computer games that interact with a user exclusively through text.

5 Assessment and Proposal for Future Work

On the one hand, using knowledge-intensive techniques, systems have been built that understand a handful of similar narratives in depth. On the other hand, using corpus-based techniques, systems have been built that understand many narratives at a shallow level. No system yet exists, however, that both (1) understands a large number of arbitrary, previously unseen narratives and (2) understands these narratives in depth. This is because there has been no major project with the goal of driving up the performance of such a system. There is no time like the present. Let's start this project. We need 20 to 30 people, and we need to address the following problems:

- **Efficiency of reasoning**

  Automated reasoning algorithms encounter combinatorial explosion when applied to narratives and knowledge bases. To address this problem, we need to develop and evaluate novel reasoning architectures that use multiple reasoning and representation methods, theory partitioning (Amir & McIlraith, 2005), parallel theorem proving, partial instantiation (Kagan, Nerode, & Subrahmanian, 1994), path planning and other domain-specific algorithms, and landmark time instead of integer and real time.

- **Effective model finding**

  Model finders produce too many models of a narrative, which renders them less useful to applications. The understanding of a narrative should consist of a small number of models. We need to develop and evaluate new model finding architectures that exploit event minimization, nonmonotonic reasoning about initial conditions, nonmonotonic reasoning about scripts, nonmonotonic reasoning about spatial layouts, and probabilistic reasoning. We need to develop efficient algorithms for model finding in the context of model preferences.

- **Representation of knowledge for narrative**

  Narratives touch on a number of realms for which well-developed representations do not yet exist. We need to develop innovative representations of areas like counterplanning, goal prioritization, metacognition, planning strategies, and reflection.

- **Acquisition of knowledge for narrative**

  Acquisition of knowledge for understanding narratives is time-consuming. To speed up acquisition, we need to develop a collaborative, workflow-based acquisition system along the lines of Open Mind Common Sense (Singh et al., 2002), in which the most important knowledge is acquired with the highest priority. We also need to work on automated acquisition and learning techniques.

- **Acquisition of annotated training data**

  For developing, training, and testing the narrative understanding system, we need to create large quantities of annotated narrative data. NarrativeML (Mani, 2013), a markup language for narrative, can be used for annotation. We need gold standard annotated narrative corpora as well as gold standard answers for tasks.

I propose the following task:

CMN Task Definition. Given a narrative text T, list of characters C, and list of objects O, answer a list of questions about C and O after every sentence of T.

I previously proposed the following list of evaluation questions for narrative understanding systems (Mueller, 2000b):

CMN Question Set. Why did P perform an action or cause a state? How did P react to an action or state? What did P expect to happen when performing an action or causing a state? What was P's goal? Did P's goal succeed or fail?

How did P feel? Why did P feel a state? How did P's emotions/feelings affect P's actions? Who/what did P like/dislike? Why did P like/dislike someone/something? How did P liking/disliking someone/something affect P's actions?

Where was someone/something? Where did someone/something go? Where did someone/something come from? How did someone/something move? Who/what was near/in/on/... someone/something? Where was someone/something in relation to someone/something else?

Who had something before an action? Who had something after an action? Who transferred something to someone?

Were the goals of P1 and P2 in concord or conflict? What were those goals? What was the outcome?

What time of day was it? How long did something take? Why did an action or state occur?

Who/what performed an action or caused a state?

What is the theme of the story?

References

Allen, J. F. (1983). Maintaining knowledge about temporal intervals. Communications of the ACM, 26(11), 832–843.

Allen, J. F. (1984). Towards a general theory of action and time. Artificial Intelligence, 23, 123–154.

Alshawi, H. (Ed.). (1992). The Core Language Engine. Cambridge, MA, USA: MIT Press.

Alvarado, S. J. (1990). Understanding editorial text: A computer model of argument comprehension. Boston, MA, USA: Kluwer.

Amir, E., & McIlraith, S. (2005). Partition-based logical reasoning for first-order and propositional theories. Artificial Intelligence, 162, 49–88.

Arbib, M. A., Conklin, E. J., & Hill, J. C. (1987). From schema theory to language. New York, NY, USA: Oxford University Press.

August, S. (1991). Ariel: An approach to understanding analogies in arguments (Tech. Rep. No. 910051). Computer Science Department, University of California, Los Angeles.

Baker, C. F., Fillmore, C. J., & Lowe, J. B. (1998). The Berkeley FrameNet project. In Proceedings of the Thirty-Sixth Annual Meeting of the Association for Computational Linguistics and Seventeenth International Conference on Computational Linguistics: Proceedings of the Conference (pp. 86–90). San Mateo, CA, USA: Morgan Kaufmann.

Bates, J., Loyall, A. B., & Reilly, W. S. (1994). An architecture for action, emotion, and social behavior. In C. Castelfranchi & E. Werner (Eds.), Artificial social systems (Vol. 830, pp. 55–68). Berlin: Springer.

Blackburn, P., & Bos, J. (2005). Representation and inference for natural language: A first course in computational semantics. Stanford, CA, USA: CSLI.

Bower, G. H. (1989). Mental models in text understanding. In A. F. Bennett & K. M. McConkey (Eds.), Cognition in individual and social contexts (pp. 129–144). Amsterdam: Elsevier.

Brachman, R. J., & Levesque, H. J. (2004). Knowledge representation and reasoning. San Francisco, CA, USA: Morgan Kaufmann.

Bringsjord, S., & Ferrucci, D. A. (2000). Artificial intelligence and literary creativity. Mahwah, NJ, USA: Lawrence Erlbaum.

Brooks, K. M. (1999). Metalinear cinematic narrative: Theory, process, and tool (Unpublished doctoral dissertation). Program in Media Arts and Sciences, School of Architecture and Planning, Massachusetts Institute of Technology, Cambridge, MA, USA.

Callaway, C. (2000). Narrative prose generation (Unpublished doctoral dissertation). Department of Computer Science, North Carolina State University, Raleigh, NC, USA.

Carbonell, J. G. (1979). Subjective understanding: Computer models of belief systems (Tech. Rep. No. YALE/DCS/tr150). New Haven, CT, USA: Computer Science Department, Yale University.

Carbonell, J. G. (1980). Towards a process model of human personality traits. Artificial Intelligence, 15, 49–74.

Carroll, G., & Charniak, E. (1991). A probabilistic analysis of marker-passing techniques for plan recognition (Tech. Rep. No. CS-91-44). Providence, RI, USA: Department of Computer Science, Brown University.

Cavazza, M., Si, M., & Zook, A. (Eds.). (2013). Intelligent narrative technologies VI: Papers from the 2013 AIIDE workshop (Vol. WS-12-14). Menlo Park, CA, USA: Association for the Advancement of Artificial Intelligence.

Chambers, N., & Jurafsky, D. (2008). Unsupervised learning of narrative event chains. In K. McKeown, J. D. Moore, S. Teufel, J. Allan, & S. Furui (Eds.), Proceedings of the Forty-Sixth Annual Meeting of the Association for Computational Linguistics (pp. 789–797). Association for Computational Linguistics.

Chambers, N., & Jurafsky, D. (2009). Unsupervised learning of narrative schemas and their participants. In K.-Y. Su, J. Su, & J. Wiebe (Eds.), Proceedings of the Forty-Seventh Annual Meeting of the Association for Computational Linguistics and the Fourth International Joint Conference on Natural Language Processing of the AFNLP (pp. 602–610). Association for Computational Linguistics.

Charniak, E. (1972). Toward a model of children's story comprehension (Tech. Rep. No. AITR-266). Cambridge, MA, USA: Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Charniak, E. (1977a). A framed PAINTING: The representation of a common sense knowledge fragment. Cognitive Science, 1, 355–394.

Charniak, E. (1977b). Ms. Malaprop, a language comprehension program. In R. Reddy (Ed.), Proceedings of the Fifth International Joint Conference on Artificial Intelligence (pp. 1–7). Los Altos, CA, USA: William Kaufmann.

Charniak, E. (1983). Passing markers: A theory of contextual influence in language comprehension. Cognitive Science, 7, 171–190.

Charniak, E. (1986). A neat theory of marker passing. In Proceedings of the Fifth National Conference on Artificial Intelligence (pp. 584–588). Los Altos, CA, USA: Morgan Kaufmann. Charniak, E., Riesbeck, C. K., & McDermott, D. V. (1980). Artificial intelligence programming. Hillsdale, NJ, USA: Lawrence Erlbaum.

Clark, H. H. (1977). Bridging. In P. N. Johnson-Laird & P. C. Wason (Eds.), Thinking: Readings in cognitive science (pp. 411–420). Cambrigde, UK: Cambridge University Press.

Clark, P. (1999). Story generation and aviation incident representation. Austin, TX, USA: Department of Computer Science, University of Texas at Austin, http://www.cs.utexas.edu/users/pclark/working_notes/.

Cox, M. T. (1996). Introspective multistrategy learning: Constructing a learning strategy under reasoning failure (Tech. Rep. No. GIT-CC-96-06). Atlanta, GA, USA: College of Computing, Georgia Institute of Technology.

Cullingford, R. E. (1978). Script application: Computer understanding of newspaper stories (Tech. Rep. No. YALE/DCS/tr116). New Haven, CT, USA: Computer Science Department, Yale University.

Dahlgren, K., McDowell, J., & Stabler, E. P., Jr. (1989). Knowledge representation for commonsense reasoning with text. Computational Linguistics, 15(3), 149–170.

Dale, R., & Reiter. (2000). Building natural language generation systems. Cambridge University Press.

Davis, E. (1990). Representations of commonsense knowledge. San Mateo, CA, USA: Morgan Kaufmann.

Davis, E. (1995). A highly expressive language of spatial constraints (Tech. Rep. No. TR1995-714). New York, NY, USA: Computer Science Department, New York University.

Dehn, N. (1981). Memory in story invention. In Proceedings of the Third Annual Conference of the Cognitive Science Society (pp. 213–215). Berkeley, CA, USA.

Dehn, N. (1989). Computer story-writing: The role of reconstructive and dynamic memory (Tech. Rep. No. YALE/DCS/tr712). New Haven, CT, USA: Computer Science Department, Yale University.

DeJong, G. F. (1979). Skimming stories in real time: An experiment in integrated understanding (Tech. Rep. No. YALE/DCS/tr158). New Haven, CT, USA: Computer Science Department, Yale University.

Dolan, C. (1989). Tensor manipulation networks: Connectionist and symbolic approaches to comprehension, learning, and planning (Tech. Rep. No. CSD-890030). Computer Science Department, University of California, Los Angeles.

Dorr, B. J. (1993). Machine translation: A view from the lexicon. Cambridge, MA, USA: MIT Press.

Duchan, J. F., Bruder, G. A., & Hewitt, L. E. (Eds.). (1995). Deixis in narrative: A Cognitive Science perspective. Hillsdale, NJ, USA: Lawrence Erlbaum.

Dyer, M. G. (1983). In-depth understanding: A computer model of integrated processing for narrative comprehension. Cambridge, MA, USA: MIT Press.

Ebbinghaus, H.-D., Flum, J., & Thomas, W. (1994). Mathematical logic (Second ed.). New York, NY, USA: Springer.

Elson, D. K. (2012). Detecting story analogies from annotations of time, action and agency. In M. A. Finlayson, P. Gervás, D. Yuret, & F. Bex (Eds.), Proceedings of the Third Workshop on Computational Models of Narrative (pp. 91–99). İstanbul, Turkey.

Emmott, C. (1997). Narrative comprehension: A discourse perspective. Oxford: Oxford University Press.

Fayzullin, M., Subrahmanian, V. S., Albanese, M., Cesarano, C., & Picariello, A. (2007). Story cre- ation from heterogeneous data sources. Multimedia Tools and Applications, 33(3), 351–377.

Feldman, J. A. (2006). From molecule to metaphor: A neural theory of language. Cambridge, MA, USA: MIT Press.

Fikes, R. E., & Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. Artificial Intelligence, 2, 189–208.

Fillmore, C. J. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), Universals in linguistic theory (pp. 1–90). New York, NY, USA: Holt, Rinehart, and Winston.

Finlayson, M. A. (Ed.). (2010). Computational models of narrative: Papers from the 2010 AAAI fall symposium (Vol. FS-10-04). Menlo Park, CA, USA: Association for the Advancement of Artificial Intelligence.

Finlayson, M. A. (Ed.). (2012). The third workshop on Computational Models of Narrative. İstanbul, Turkey.

Finlayson, M. A., Fisseni, B., Löwe, B., & Meister, J. C. (2013). Proceedings of the 2013 Workshop on Computational Models of Narrative. Hamburg, Germany.

Finlayson, M. A., Richards, W., & Winston, P. H. (2010). Computational models of narrative: Review of a workshop. AI Magazine, 31(2), 97–100.

Fisher, M., Gabbay, D. M., & Vila, L. (Eds.). (2005). Handbook of temporal reasoning in artificial intelligence. Amsterdam: Elsevier.

Fisseni, B., Kurji, A., & Löwe, B. (2014). Annotating with Propp's Morphology of the Folk-tale: Reproducibility and Trainability. Literary and Linguistic Computing, 29, 488–510.

Frank, S. L., Koppen, M., Noordman, L. G. M., & Vonk, W. (2003). Modeling knowledge-based inferences in story comprehension. Cognitive Science, 27, 875–910.

Gabbay, D. M., & Ohlbach, H. J. (Eds.). (1994). Temporal logic. Berlin: Springer.

Gerard, S., & Sansonnet, J.-P. (2000). A spatio-temporal model for the representation of situations described in narrative texts. In D. Christodoulakis (Ed.), Natural language processing (Vol. 1835, pp. 176–184). Berlin: Springer.

Gervás, P. (2012). Story generator algorithms. In P. Hühn, J. C. Meister, W. Schernus, J. Pier, W. Schmid, & J. Schönert (Eds.), The living handbook of narratology. Hamburg: Hamburg University Press. Retrieved from http://www.lhn.uni-hamburg.de/article/story-generator-algorithms

Gervás, P. (2013). Propp's morphology of the folk tale as a grammar for generation. In M. A. Fin- layson, B. Fisseni, B. Löwe, & J. C. Meister (Eds.), Proceedings of the 2013 Workshop on Computational Models of Narrative. Hamburg, Germany.

Gervás, P., Díaz-Agudo, B., Peinado, F., & Hervás, R. (2005). Story plot generation based on CBR. Knowledge-Based Systems, 18(4–5), 235–242.

Ghallab, M., Nau, D., & Traverso, P. (2004). Automated planning: Theory and practice. San Francisco, CA, USA: Morgan Kaufmann.

Gildea, D., & Jurafsky, D. (2002). Automatic labeling of semantic roles. Computational Linguistics, 28(3), 245–288.

Golden, R. M., & Rumelhart, D. E. (1993). A parallel distributed processing model of story comprehension and recall. Discourse Processes, 16, 203–237.

Goldman, R. P. (1990). A probabilistic approach to language understanding (Tech. Rep. No. CS-90-34). Providence, RI, USA: Department of Computer Science, Brown University.

Goldman, S. R., Graesser, A. C., & van den Broek, P. (Eds.). (1999). Narrative comprehension, causality, and coherence: Essays in honor of Tom Trabasso. Mahwah, NJ, USA: Lawrence Erlbaum.

Gordon, A. S. (1999). The design of knowledge-rich browsing interfaces for retrieval in digital libraries (Unpublished doctoral dissertation). Northwestern University, Evanston, IL, USA. Gordon, A. S. (2000). The Commonsense Activity Resource Page: Cataloging the activities that everyone knows about. Institute for Creative Technologies, University of Southern California. Retrieved from http://people.ict.usc.edu/~gordon/carp.html

Goyal, A., Riloff, E., & Daumé, III, H. (2010). Automatically producing plot unit representations for narrative text. In Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing (pp. 77–86). Association for Computational Linguistics.

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

Heider, F. (1958). The psychology of interpersonal relations. Hillsdale, NJ, USA: Lawrence Erlbaum.

Hobbs, J. R., Stickel, M. E., Appelt, D. E., & Martin, P. (1993). Interpretation as abduction. Artificial Intelligence, 63, 69–142.

Jackendoff, R. (1972). Semantic interpretation in generative grammar. Cambridge, MA, USA: MIT Press.

Jackendoff, R. (1990). Semantic structures. Cambridge, MA, USA: MIT Press.

Jhala, A. H., Riedl, M. O., & Roberts, D. L. (Eds.). (2010). INT3 '10: Proceedings of the intelligent narrative technologies III workshop. New York, NY, USA: ACM.

Johnson-Laird, P. N. (1983). Mental models: Toward a cognitive science of language, inference, and consciousness. Cambridge, MA, USA: Harvard University Press.

Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. Psychological Review, 87(4), 329–354.

Kagan, V., Nerode, A., & Subrahmanian, V. S. (1994). Computing definite logic programs by partial instantiation. Annals of Pure and Applied Logic, 67, 161–182.

Kass, A. M. (1990). Developing creative hypotheses by adapting explanations (Unpublished doctoral dissertation). Computer Science Department, Yale University, New Haven, CT, USA. Kautz, H. A. (1991). A formal theory of plan recognition and its implementation. In J. F. Allen, H. A. Kautz, R. N. Pelavin, & J. D. Tenenberg (Eds.), Reasoning about plans (pp. 69–125). San Mateo, CA, USA: Morgan Kaufmann.

Kingsbury, P., & Palmer, M. (2002). From TreeBank to PropBank. In Proceedings of the Third International Conference on Language Resources and Evaluation. Paris: Evaluations and Language Resources Distribution Agency.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. Psychological Review, 95(163–182).

Kintsch, W. (1998). Comprehension: A paradigm for cognition. Cambrigde, UK: Cambridge University Press.

Kintsch, W., & van Dijk, T. A. (1978). Toward a model of text comprehension and production. Psychological Review, 85, 363–394.

Klein, S., Aeschlimann, J. F., Balsiger, D. F., Converse, S. L., Court, C., Foster, M., Lao, R., Oakley, J. D., & Smith, J. (1973). Automatic novel writing: A status report (Tech. Rep. No. 186). Department of Computer Sciences, University of Wisconsin-Madison.

Kolodner, J. L. (1984). Retrieval and organizational strategies in conceptual memory: A computer model. Hillsdale, NJ, USA: Lawrence Erlbaum.

Kolodner, J. L. (1993). Case-based reasoning. San Mateo, CA, USA: Morgan Kaufmann.

Kuflik, T., Nissan, E., & Puni, G. (1991). Finding excuses with ALIBI: Alternative plans that are deontically more defensible. Computers and Artificial Intelligence, 10, 297-325.

Kuipers, B. (2000). The Spatial Semantic Hierarchy. Artificial Intelligence, 119, 191–233. Lakoff, G., & Johnson, M. (1980). Metaphors we live by. Chicago: University of Chicago Press. Lakoff, G., & Narayanan, S. (2010). Toward a computational model of narrative. In M. A. Finlayson, P. Gervás, E. T. Mueller, S. Narayanan, & P. H. Winston (Eds.), Computational models of narrative: Papers from the AAAI Fall Symposium (pp. 21–28). Menlo Park, CA, USA: AAAI Press.

Lang, R. R. (1997). A formal model for simple narratives (Unpublished doctoral dissertation). Tulane University, New Orleans, LA, USA.

Lange, T. E., & Dyer, M. G. (1989). High-level inferencing in a connectionist network. Connection Science, 1, 181–217.

Lange, T. E., & Wharton, C. M. (1992). Remind: Retrieval from episodic memory by inferencing and disambiguation (Tech. Rep. No. 920047). Computer Science Department, University of California, Los Angeles.

Langston, M. C., Trabasso, T., & Magliano, J. P. (1999). A connectionist model of narrative comprehension. In A. Ram & K. Moorman (Eds.), Understanding language understanding: Computational models of reading (pp. 181–226). Cambridge, MA, USA: MIT Press.

Laurel, B. K. (1986). Toward the design of a computer-based interactive fantasy system (Unpublished doctoral dissertation). Ohio State University, Columbus, OH, USA.

Leake, D. B. (1992). Evaluating explanations: A content theory. Hillsdale, NJ, USA: Lawrence Erlbaum.

Leake, D. B. (Ed.). (1996). Case-based reasoning: Experiences, lessons, & future directions. Cambridge, MA, USA: MIT Press.

Lebowitz, M. (1980). Generalization and memory in an integrated understanding system (Tech. Rep. No. YALE/DCS/tr186). New Haven, CT, USA: Computer Science Department, Yale University.

Lebowitz, M. (1984). Creating characters in a story-telling universe. Poetics, 13, 171–194.

Lebowitz, M. (1985). Story-telling as planning and learning. Poetics, 14, 483–502.

Lehnert, W. G. (1978). The process of question answering: A computer simulation of cognition. Hillsdale, NJ, USA: Lawrence Erlbaum.

Lehnert, W. G. (1982). Plot units: A narrative summarization strategy. In W. G. Lehnert & M. H. Ringle (Eds.), Strategies for natural language processing (pp. 375–412). Hillsdale, NJ, USA: Lawrence Erlbaum.

Lehnert, W. G. (1994). Cognition, computers, and car bombs: How Yale prepared me for the 1990s. In R. C. Schank & E. Langer (Eds.), Beliefs, reasoning, and decision making: Psycho-logic in honor of Bob Abelson (pp. 143–173). Hillsdale, NJ, USA: Lawrence Erlbaum.

Lenat, D. B., & Guha, R. V. (1990). Building large knowledge-based systems: Representation and inference in the Cyc project. Reading, MA, USA: Addison-Wesley.

Library of Congress. (1995). Thesaurus for graphic materials. http://www.loc.gov/rr/print/tgm1/.

Liu, H., & Singh, P. (2002). MAKEBELIEVE: Using commonsense knowledge to generate stories. In Proceedings of the Eighteenth National Conference on Artificial Intelligence and Fourteenth Conference on Innovative Applications of Artificial Intelligence (pp. 957–958). Menlo Park, CA, USA: AAAI Press.

Louchart, S., Mehta, M., & Roberts, D. L. (Eds.). (2009). Intelligent narrative technologies II: Papers from the 2009 AAAI spring symposium (Vol. SS-09-06). Menlo Park, CA, USA: Association for the Advancement of Artificial Intelligence.

Löwe, B., Pacuit, E., & Saraf, S. (2009). Identifying the structure of a narrative via an agent-based logic of preferences and beliefs: Formalizations of episodes from CSI: Crime Scene Investigation. In M. Duvigneau & D. Moldt (Eds.), Proceedings of the Fifth International Workshop on Modelling of Objects, Components, and Agents (pp. 45–63). Hamburg, Germany.

Loyall, A. B. (1997). Believable agents: Building interactive personalities (Tech. Rep. No. CMU-CS-97-123). Pittsburgh, PA, USA: School of Computer Science, Carnegie Mellon University.

Magerko, B. S., & Riedl, M. O. (Eds.). (2007). Intelligent narrative technologies: Papers from the 2007 AAAI fall symposium (Vol. FS-07-05). Menlo Park, CA, USA: Association for the Advancement of Artificial Intelligence.

Mahesh, K. (1995). Syntax-semantics interaction in sentence understanding (Tech. Rep. No. GIT-CC-95-10). Atlanta, GA, USA: College of Computing, Georgia Institute of Technology.

Manandhar, S., & Yuret, D. (2013). Semeval-2013: Semantic evaluation exercises. http://www.cs.york.ac.uk/semeval-2013/.

Mani, I. (2013). Computational modeling of narrative. San Rafael, CA, USA: Morgan & Claypool. Manning, C. D., & Schütze, H. (1999). Foundations of statistical natural language processing. Cambridge, MA, USA: MIT Press.

Mateas, M., & Stern, A. (2003). Façade: An experiment in building a fully-realized interactive drama. In Game Developers Conference.

McCarthy, J. (1959). Programs with common sense. In Mechanisation of Thought Processes: Proceedings of a Symposium held at the National Physical Laboratory on 24th, 25th, 26th and 27th November 1958 (Vol. 1, pp. 75–91). London: Her Majesty's Stationery Office.

McCarthy, J. (1990). An example for natural language understanding and the AI problems it raises. In V. Lifschitz (Ed.), Formalizing common sense: Papers by John McCarthy (pp. 70–76). Norwood, NJ, USA: Ablex.

McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (Eds.), Machine intelligence 4 (pp. 463–502). Edinburgh, Scotland: Edinburgh University Press.

McCarthy, J., Minsky, M., Sloman, A., Gong, L., Lau, T., Morgenstern, L., Mueller, E. T., Riecken, D., Singh, M., & Singh, P. (2002). An architecture of diversity for commonsense reasoning. IBM Systems Journal, 41(3), 530–539. McClelland, J. L., Rumelhart, D. E., & PDP Research Group (Eds.). (1986). Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 2: Psychological and Biological Models). Cambridge, MA, USA: MIT Press.

McCord, M. C., Murdock, J. W., & Boguraev, B. K. (2012). Deep parsing in Watson. IBM Journal of Research and Development, 56(3/4), 3:1–3:15.

McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. Journal of Experimental Psychology: Learning, Memory, and Cognition, 12(1), 82–91.

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

Meehan, J. (1976). The metanovel: Writing stories by computer (Tech. Rep. No. YALE/DCS/tr074). New Haven, CT, USA: Computer Science Department, Yale University.

Michael, L. (2010). Computability of narrative. In M. A. Finlayson, P. Gervás, E. T. Mueller, S. Narayanan, & P. H. Winston (Eds.), Computational models of narrative: Papers from the AAAI Fall Symposium (pp. 47–54). Menlo Park, CA, USA: AAAI Press.

Michael, L. (2012). Similarity of narratives. In M. A. Finlayson, P. Gervás, D. Yuret, & F. Bex (Eds.), Proceedings of the Third Workshop on Computational Models of Narrative (pp. 105–113). İstanbul, Turkey.

Michael, L. (2013). Story understanding ... calculemus! In A. Kakas (Ed.), Proceedings of the Eleventh International Symposium on Logical Formalizations of Commonsense Reasoning. Miikkulainen, R. (1993). Subsymbolic natural language processing. Cambridge, MA, USA: MIT Press.

Miller, R., & Shanahan, M. (2002). Some alternative formulations of the event calculus. In A. C. Kakas & F. Sadri (Eds.), Computational logic: Logic programming and beyond: Essays in honour of Robert A. Kowalski, part II (Vol. 2408, pp. 452–490). Berlin: Springer.

Minsky, M. (1961). Steps toward artificial intelligence. Proceedings of the IRE, 49(1), 8–30. Minsky, M. (1974). A framework for representing knowledge. A. I. Memo 306, Cambridge, MA, USA: Artificial Intelligence Laboratory, Massachusetts Institute of Technology. Minsky, M. (1986). The society of mind. New York, NY, USA: Simon & Schuster.

Montfort, N. (2003). Twisty little passages: An approach to interactive fiction. Cambridge, MA, USA: MIT Press.

Moorman, K. (1997). A functional theory of creative reading: Process, knowledge, and evaluation (Tech. Rep. Nos. GIT–CC–94/01). Atlanta, GA, USA: College of Computing, Georgia Institute of Technology.

Morrow, D. (1994). Spatial models created from text. In H. van Oostendorp & R. A. Zwaan (Eds.), Naturalistic text comprehension (pp. 57–78). Norwood, NJ, USA: Ablex.

Mross, E. F., & Roberts, J. O. (1992). The construction-integration model: A program and manual (Tech. Rep. No. 92-14). Boulder, CO, USA: Institute of Cognitive Science, University of Colorado at Boulder.

Mueller, E. T. (1990). Daydreaming in humans and machines: A computer model of the stream of thought. Norwood, NJ, USA: Ablex.

Mueller, E. T. (1998). Natural language processing with ThoughtTreasure. New York, NY, USA: Signiform.

Mueller, E. T. (2000a). A database and lexicon of scripts for ThoughtTreasure. Computing Research Repository, cs.AI/0003004.

Mueller, E. T. (2000b). Prospects for in-depth story understanding by computer. Computing Research Repository, cs.AI/0003003.

Mueller, E. T. (2002). Story understanding. In L. Nadel (Ed.), Encyclopedia of cognitive science (Vol. 4, pp. 238–246). London: Nature Publishing Group.

Mueller, E. T. (2003). Story understanding through multi-representation model construction. In G. Hirst & S. Nirenburg (Eds.), Text Meaning: Proceedings of the HLT-NAACL 2003 Workshop (pp. 46–53). East Stroudsburg, PA, USA: Association for Computational Linguistics.

Mueller, E. T. (2004). Understanding script-based stories using commonsense reasoning. Cognitive Systems Research, 5(4), 307–340.

Mueller, E. T. (2006). Commonsense reasoning. San Francisco, CA, USA/Amsterdam: Morgan Kaufmann/Elsevier.

Mueller, E. T. (2007a). Modelling space and time in narratives about restaurants. Literary and Linguistic Computing, 22(1), 67–84.

Mueller, E. T. (2007b). Understanding goal-based stories through model finding and planning. In B. S. Magerko & M. O. Riedl (Eds.), Intelligent narrative technologies: Papers from the AAAI Fall Symposium (pp. 95–101). Menlo Park, CA, USA: AAAI Press.

Mulkar-Mehta, R. (2000). Mini-TACITUS. http://www.rutumulkar.com/tacitus.html.

Nackoul, D. D. (2010). Text to text: Plot unit searches generated from English (Unpublished master's thesis). Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA.

Narayanan, S. S. (1997). Knowledge-based action representations for metaphor and aspect (KARMA) (Unpublished doctoral dissertation). University of California, Berkeley, Berkeley, CA, USA.

Newell, A., & Simon, H. A. (1961). GPS, a program that simulates human thought. In H. Billing (Ed.), Lernende automaten (pp. 109–124). Munich: Oldenbourgh.

Norvig, P. (1987). A unified theory of inference for text understanding (Tech. Rep. No. CSD-87-339). Berkeley, CA, USA: Computer Science Division, University of California, Berkeley.

Nossum, R. (2003). A contextual approach to the logic of fiction. In P. Blackburn, C. Ghidini, R. M. Turner, & F. Giunchiglia (Eds.), Modeling and using context (Vol. 2680, pp. 233–244). Berlin: Springer.

O'Rorke, P., & Ortony, A. (1994). Explaining emotions. Cognitive Science, 18, 283–323. Ortony, A., Clore, G. L., & Collins, A. M. (1988). The cognitive structure of emotions. Cambridge, UK: Cambridge University Press.

Owens, C. C. (1990). Indexing and retrieving abstract planning knowledge (Unpublished doctoral dissertation). Computer Science Department, Yale University, New Haven, CT, USA.

Palmer, M., Gildea, D., & Kingsbury, P. (2005). The Proposition Bank: An annotated corpus of semantic roles. Computational Linguistics, 31(1), 71–105.

Palmer, M., Passonneau, R. J., Weir, C., & Finin, T. W. (1993). The KERNEL text understanding system. Artificial Intelligence, 63(1–2), 17–68.

Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval, 2(1–2), 1–135.

Pérez y Pérez, R., & Sharples, M. (2001). MEXICA: A computer model of a cognitive account of creative writing. Journal of Experimental and Theoretical Artificial Intelligence, 13(2), 119–139.

Phillips, M. A., & Huntley, C. (1993). Dramatica: A theory of story. http://dramatica.com/theory/book.

Prince, G. (1973). A grammar of stories. The Hague: Mouton.

Prince, G. (1982). Narratology: The form and functioning of narrative. Berlin: Mouton. Propp, V. (1968). Morphology of the folktale. Austin, TX, USA: University of Texas Press.

Punyakanok, V., Roth, D., & Yih, W.-t. (2008). The importance of syntactic parsing and inference in semantic role labeling. Computational Linguistics, 34(2), 257–287.

Racter. (1984). The policeman's beard is half constructed. Warner Books.

Ram, A. (1989). Question-driven understanding: An integrated theory of story understanding, memory, and learning (Tech. Rep. No. YALE/DCS/tr710). New Haven, CT, USA: Computer Science Department, Yale University.

Ram, A., & Moorman, K. (Eds.). (1999). Understanding language understanding: Computational models of reading. Cambridge, MA, USA: MIT Press.

Randell, D. A., Cui, Z., & Cohn, A. G. (1992). A spatial logic based on regions and connection. In B. Nebel, C. Rich, & W. R. Swartout (Eds.), Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (pp. 165–176). San Mateo, CA, USA: Morgan Kaufmann.

Reeves, J. (1991). Computational morality: A process model of belief conflict and resolution for story understanding (Tech. Rep. No. CSD-910017). Computer Science Department, University of California, Los Angeles.

Reilly, W. S. N. (1996). Believable social and emotional agents (Tech. Rep. No. CMU-CS-96-138). Pittsburgh, PA, USA: School of Computer Science, Carnegie Mellon University.

Reisig, W. (1985). Petri nets: An introduction. Berlin: Springer.

Reiter, R. (2001). Knowledge in action: Logical foundations for specifying and implementing dynamical systems. Cambridge, MA, USA: MIT Press.

Richards, W., Finlayson, M. A., & Winston, P. H. (2009). Advancing computational models of narrative (Tech. Rep. No. MIT-CSAIL-TR-2009-063). Cambridge, MA, USA: Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

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

Riesbeck, C. K., & Schank, R. C. (1989). Inside case-based reasoning. Hillsdale, NJ, USA: Lawrence Erlbaum.

Riloff, E. (1999). Information extraction as a stepping stone toward story understanding. In A. Ram & K. Moorman (Eds.), Understanding language understanding: Computational models of reading (pp. 435–460). Cambridge, MA, USA: MIT Press.

Rimmon-Kenan, S. (2002). Narrative fiction: Contemporary poetics (2nd ed.). London: Routledge.

Rosenberg, S. T. (1977). Frame-based text processing. A. I. Memo 431, Cambridge, MA, USA: Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Rumelhart, D. E. (1975). Notes on a schema for stories. In D. G. Bobrow & A. M. Collins (Eds.), Representation and understanding: Studies in cognitive science (pp. 211–236). New York, NY, USA: Academic Press.

Rumelhart, D. E., McClelland, J. L., & PDP Research Group (Eds.). (1986). Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 1: Foundations). Cambridge, MA, USA: MIT Press.

Ryan, M.-L. (1991). Possible worlds, artificial intelligence and narrative theory. Bloomington, IN, USA: Indiana University Press.

Sack, W. (1992). Micro-Talespin: A story generator. Cambridge, MA, USA: Media Laboratory, Massachusetts Institute of Technology, http://eliterature.org/images/microtalespin.txt.

Schank, R. C. (1972). Conceptual dependency: A theory of natural language understanding. Cognitive Psychology, 3(4), 552–631.

Schank, R. C. (1982). Dynamic memory. Cambridge, UK: Cambridge University Press.

Schank, R. C. (1986). Explanation patterns: Understanding mechanically and creatively.
Hillsdale, NJ, USA: Lawrence Erlbaum.

Schank, R. C., & Abelson, R. P. (1977). Scripts, plans, goals, and understanding: An inquiry into human knowledge structures. Hillsdale, NJ, USA: Lawrence Erlbaum.

Schank, R. C., Goldman, N., Rieger, C., & Riesbeck, C. K. (1973). MARGIE: Memory, analysis, response generation and inference in English. In Advance Papers of the

Third International Joint Conference on Artificial Intelligence (pp. 255–261). Stanford, CA, USA.

Schank, R. C., Goldman, N., Rieger, C., & Riesbeck, C. K. (1975). Inference and paraphrase by computer. Journal of the ACM, 22(3), 309–328.

Schank, R. C., Kass, A., & Riesbeck, C. K. (Eds.). (1994). Inside case-based explanation. Hillsdale, NJ, USA: Lawrence Erlbaum.

Schank, R. C., & Riesbeck, C. K. (Eds.). (1981). Inside computer understanding: Five programs plus miniatures. Hillsdale, NJ, USA: Lawrence Erlbaum.

Schubert, L. K., & Hwang, C. H. (1989). An episodic knowledge representation for narrative texts. In R. J. Brachman, H. J. Levesque, & R. Reiter (Eds.), Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (pp. 444–458). San Mateo, CA, USA: Morgan Kaufmann.

Schubert, L. K., & Hwang, C. H. (2000). Episodic Logic meets Little Red Riding Hood: A comprehensive, natural representation for language understanding. In L. M. Iwańska & S. C. Shapiro (Eds.), Natural language processing and knowledge representation (pp. 111–174). Cambridge, MA, USA: MIT Press.

Selfridge, O. G. (1959). Pandemonium: A paradigm for learning. In Mechanization of Thought Processes (Vol. 1, pp. 511–529). London.

Shapiro, S. C., & Rapaport, W. J. (1995). An introduction to a computational reader of narratives. In J. F. Duchan, G. A. Bruder, & L. E. Hewitt (Eds.), Deixis in narrative: A Cognitive Science perspective (pp. 79–105). Hillsdale, NJ, USA: Lawrence Erlbaum.

Singh, P., Lin, T., Mueller, E. T., Lim, G., Perkins, T., & Zhu, W. L. (2002). Open Mind Common Sense: Knowledge acquisition from the general public. In R. Meersman & Z. Tari (Eds.), On the move to meaningful internet systems 2002: DOA/COOPIS/ODBASE 2002 (Vol. 2519, pp. 1223–1237). Berlin: Springer.

Slobin, D. I. (2003). The many ways to search for a frog: Linguistic typology and the expression of motion events. In S. Strömqvist & L. Verhoeven (Eds.), Relating events in narrative: Typological and contextual perspectives (pp. 219–257). Mahwah, NJ, USA: Lawrence Erlbaum.

Smith, S., & Bates, J. (1989). Towards a theory of narrative for interactive fiction (Tech. Rep. No. CMU-CS-89-121). Pittsburgh, PA, USA: School of Computer Science, Carnegie Mellon University.

Smith, T. C., & Witten, I. H. (1991). A planning mechanism for generating story text. Literary and Linguistic Computing, 6(2), 119–126.

St. John, M. F. (1992). The Story Gestalt: A model of knowledge-intensive processes in text comprehension. Cognitive Science, 16, 271-306.

Szilas, N., & Rety, J.-H. (2004). Minimal structure for stories. In B. Barry & K. M. Brooks (Eds.), Proceedings of the First ACM Workshop on Story Representation, Mechanism & Context (pp. 25–32). New York, NY, USA: ACM.

ter Meulen, A. G. B. (1995). Representing time in natural language. Cambridge, MA, USA: MIT Press.

TimeML Working Group. (2005). TimeML 1.2.1 specifications. http://timeml.org/site/ index.html.

Tomai, E., Elson, D. K., & Rowe, J. (Eds.). (2011). Intelligent narrative technologies IV: Papers from the 2011 AIIDE workshop (Vol. WS-11-18). Menlo Park, CA, USA: Association for the Advancement of Artificial Intelligence.

Turner, S. R. (1994). The creative process. Hillsdale, NJ, USA: Lawrence Erlbaum.

van Dijk, T. A., & Kintsch, W. (1983). Strategies of discourse comprehension. Orlando, FL, USA: Academic Press.

van Harmelen, F., Lifschitz, V., & Porter, B. (Eds.). (2008). Handbook of knowledge representation. Amsterdam: Elsevier.

Viegas, E. (1999). The manifesto of large-scale semantic lexicon acquisition. Traitement automatique des langues, 40(1), 27–53.

Ware, S. G., Zhu, J., & Hodhod, R. (Eds.). (2012). Intelligent narrative technologies V: Papers from the 2012 AIIDE workshop (Vol. WS-12-14). Menlo Park, CA, USA: Association for the Advancement of Artificial Intelligence.

Weyhrauch, P. (1997). Guiding interactive drama (Tech. Rep. No. CMU-CS-97-109). Pittsburgh, PA, USA: School of Computer Science, Carnegie Mellon University.

Wilensky, R. (1978). Understanding goal-based stories (Tech. Rep. No. YALE/DCS/tr140). New Haven, CT, USA: Computer Science Department, Yale University.

Wilensky, R. (1982). Points: A theory of the structure of stories in memory. In W. G. Lehnert & M. H. Ringle (Eds.), Strategies for natural language processing (pp. 345–374). Hillsdale, NJ, USA: Lawrence Erlbaum.

Wilensky, R. (1983). Planning and understanding: A computational approach to human reasoning. Reading, MA, USA: Addison-Wesley.

Winograd, T. (1972). Understanding natural language. New York, NY, USA: Academic Press.

Wu, D. (1992). Automatic inference: A probabilistic basis for natural language interpretation (Tech. Rep. No. CSD-92-692). Berkeley, CA, USA: Computer Science Division, University of California, Berkeley.

Yazdani, M. (1983). Generating events in a fictional world of stories (Tech. Rep. No. R-113). Exeter: Department of Computer Science, University of Exeter.

Zarri, G. P. (1996). NKRL, a knowledge representation language for narrative natural language processing. In Proceedings of the Sixteenth International Conference on Computational Linguistics (pp. 1032–1035). Copenhagen.

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language and memory. Psychological Bulletin, 123(2), 162–185.