Using Narrative Function to Extract Qualitative Information from Natural Language Texts

Clifton McFate, Kenneth D. Forbus, Thomas R. Hinrichs
Qualitative Reasoning Group, Northwestern University
2133 Sheridan Road, Evanston, IL, 60208, USA
mcfateclifton79@gmail.com, {forbus, t-hinrichs}@northwestern.edu

Abstract

The naturalness of qualitative reasoning suggests that qualitative representations might be an important component of the semantics of natural language. Prior work showed that frame-based representations of qualitative process theory constructs could indeed be extracted from natural language texts. That technique relied on the parser recognizing specific syntactic constructions, which had limited coverage. This paper describes a new approach, using narrative function to represent the higher-order relationships between the constituents of a sentence and between sentences in a discourse. We outline how narrative function combined with query-driven abduction enables the same kinds of information to be extracted from natural language texts. Moreover, we also show how the same technique can be used to extract type-level qualitative representations from text, and used to improve performance in playing a strategy game.

Introduction

Qualitative representations were developed in part to serve as a formal language for expressing the contents of human mental models about continuous systems. Since such knowledge is often expressed in natural language, it makes sense to explore how qualitative representations might be used in natural language semantics. Kuehne (2004) showed that the constructs of qualitative process (QP) theory (Forbus, 1984) could be recast in a frame-based representation, compatible with the frame semantic representations used in Fillmore et al.'s (2001) FrameNet. In frame semantics, frames represent conceptual structures that are connected to lexical items through their slots. For example, the notion of qualitative proportionality is captured by an Indirect Influence frame, which includes slots for *constrainer* (the causally antecedent quantity), *constrained* (the consequent quantity), and *sign* (the direction of change). Kuehne (2004) identified a set of phrasal patterns that could be identified by syntactic parsers and used to extract QP information from natural language texts, e.g. "As the temperature of the steam rises, the pressure of the boiler rises." would lead to the addition of a qualitative proportionality with the constrainer quantity being the steam's temperature, the constrained quantity being the boiler pressure, and the sign being positive. For each representational element in QP theory (i.e. quantities, ordinals, influences, and processes), Kuehne identified a set of syntactic patterns that could be used to extract them from text. The syntactic patterns were encoded into the grammar of the parser, which is capable of using semantic constituents (e.g. sub-elements identified as quantities) in its rules. The extracted knowledge was further transformed by antecedent rules to construct QP frame representations. When trying to scale this approach up for learning by reading, we discovered two limitations. First, the use of syntactic patterns significantly limited coverage. Second, using separate mechanisms for recognizing QP information seemed inelegant: Integrating these representations into a broader, more robust framework seemed necessary.

This paper describes a different approach, based on narrative function, for extracting QP information from text. We start by explaining the idea of narrative function and the key properties of the natural language understanding system used. Then we show how QP frames can be constructed by deriving these narrative functions, and that this approach already captures the full range of examples handled previously, and provides respectable performance on simplified English text from a science book. Moreover, we show how narrative function can be used to extract type-level influences (Hinrichs & Forbus, 2012) from natural language, and that such information can significantly improve the performance of a system playing a strategy game. We close with related and future work.

Narrative Function and Abduction

When people read, they look to see how what they are reading fits together. At the beginning of a story, characters are introduced, and expectations raised about possible events that might occur. If a fable involves a fox and a goose meeting on a riverbank, for example, one possible outcome of that meeting is that mayhem ensues. Narrative function (Labov & Waletzky, 1966; Barthes 1977; Trabasso et al. 1984) provides a level of representation that ties the contents of specific sentences to the ongoing discourse. Introducing a character is a narrative function, as is introducing an event and raising expectations about possible outcomes of that event.

Tomai & Forbus (2009) showed that narrative functions could be used in understanding natural language texts such as fables and the materials found in psychological studies of social cognition and moral decision-making. Since qualitative information is part of what is conveyed in language, e.g. explanations of continuous systems, such as found in textbooks, it stands to reason that such information needs to be linked into the general-purpose representations for understanding the intended purpose of a sentence within a discourse. Thus it makes sense to expand the range of narrative functions to include detecting the introduction of QP information. This section summarizes how narrative function detection works in our natural language system, setting the stage for the new narrative functions introduced in the next section.

The Explanation Agent Natural Language Understanding System (EA NLU; Tomai & Forbus, 2009) uses a syntactic parser (Allen, 1994) and lexical information from COMLEX (Grishem et al. 1993) for syntactic processing. It also uses lexical and semantic representations from ResearchCyc[1], extended with an implementation of Discourse Representation Theory (Kamp & Reyle, 1993) that uses Cyc microtheories to handle contexts.

Like other NLU systems, EA NLU introduces choice sets to represent ambiguities. Choice sets are introduced when there are multiple meanings of a word, or multiple parses. Consider for example this discourse fragment:

"The temperature of the boiler is increasing."

One ambiguity in the sentence is the meaning of the word *temperature*. In this context, it clearly refers to a unit of measure, but another potential meaning could be a *fever*, as in "*The child has a temperature.*" These ambiguities are preserved for disambiguation as a set of disjunctive choices. Abduction has long been used in semantic interpretation (Hobbs 2004), but it tends to be intractable as the number of statements grows. Tomai & Forbus (2009) showed that by using top-down expectations, e.g. narrative functions, many potential choices were irrelevant, thus greatly reducing the complexity of abduction. The process is driven incrementally, by first finding appropriate queries to ask, given the current logical environment, i.e.

`(queryForInterpretation ?o ?q)`

returns bindings for ?q that are queries to be made. ?o is an integer that specifies the ordering of queries, i.e. a query at level *n* can assume that all queries at level *n-1* have already been performed. Thus, for example, the rules searching for influences can be assured that information about quantities will already have been found.

The abduction mechanism is tuned for specific tasks and contexts in two ways. First, all analyses are done with respect to a logical environment, defined by a current microtheory and all of the microtheories it inherits from. This includes microtheories that specify what questions make sense for that task via `queryForInterpretation` statements. Second, the algorithm retrieves declarative advice from the logical environment as to what sorts of interpretation are preferred. For example, interpretations which include QP information are preferred, which biases the system toward interpretations that produce this sort of information.

The queries concerning narrative functions take the following form:

`(narrativeFunction ?PE ?C ?T)`

where `?PE` is a *presentation event*, i.e. the narrative-level event being described, `?C` is the content of that event, and `?T` is the type of narrative event. A sentence can give rise to multiple narrative functions, so presentation events are represented via non-atomic terms as follows:

`(PresentationEventFn <sentence ID> ?eventID)`

where `<sentence ID>` is an identifier for the sentence being processed, a meta-variable automatically substituted into each query, and `?eventID` is a unique identifier constructed by whatever rule introduces the presentation event. In the case of QP language interpretation, the content of events are particular types of QP frames from the ontology outlined below.

Finding QP frames via Abduction

Next we describe the narrative functions for QP frames that we have developed, and summarize some important properties of the rules that derive them from the natural language analysis of texts.

[1] www.cyc.com/platform/researchcyc

For each QP Frame type, we introduce a category of narrative function (see Table 1). Each frame type has a set of frame elements (aka slots). For example, a quantity frame has the following frame elements:

- `entity`: The entity that the quantity is part of.
- `quantityType`: The type of continuous property
- `value`: A numerical value (optional)
- `unit`: Units associated with a numerical value (optional)
- `signOfDerivative`: -1, 0, or 1 (optional)

Frame	Narrative Function
Quantities	IntroductionOfQuantityEvent
Topological Constraints	IntroductionOfTopologyConstraint
Derivative Sign	IntroductionOfDsInformation
Ordinals	IntroductionOfOrdinalEvent
Indirect Influence	IntroductionOfQPropEvent
Direct Influence	IntroductionOfDirectInfluenceEvent
Quantity Transfer	IntroductionOfQuantityTransferFrame
Process Frame	IntroductionOfProcess
Process Roles	IntroducesProcessRole

Table 1: QP Narrative Functions

For example, Figure 1 shows the direct influence frame built for a sentence from a science book (part of the corpus used below). The set of frame elements used is the same as Kuehne (2004), with one extension. We created a frame type for describing topological constraints on a system such as connections, interruptions, and paths. For example, in the sentence "*Water flows through a pipe.*" the path of the flow, the pipe, would be represented in a topological constraint frame. This separation was necessary as topological constraints on physical systems can frequently appear in text separated from the physical event that they constrain, e.g. *"Cylinder A1 is connected to Cylinder A2 by a pipe. Water flows from Cylinder A1 to A2."*

Solutions to narrative function queries are found via Horn clause rules[2]. These rules analyze the predicate calculus statements produced by the parser, including lexical, syntactic, and semantic information. For example, a common indicator of a quantity is a phrase like "temperature in the reactor". The prepositional phrase involving "in" leads to the parser producing an `in-UnderspecifiedContainer` statement. This is a high-level Cyc predicate that covers a large space of more specific possibilities. When the phrase that is being modified is a type of continuous quantity (here, temperature), a rule looking for this combination hypothesizes a quantity frame whose `entity` is the discourse variable for the noun in the prepositional phrase and whose `quantityType` is the kind of continuous parameter being modified.

[2] Unlike Prolog, all solutions are found, there is no notion of cut, and dependency information is recorded.

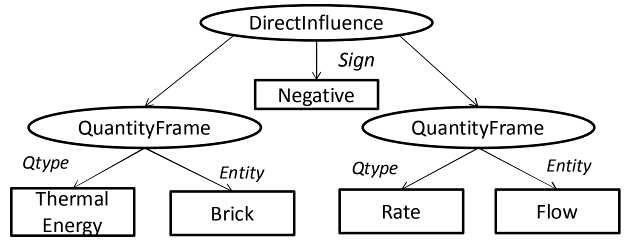

Figure 1: Example: "Heat flows from the brick"

Other rules require more type-level reasoning. For example, phrases that mention a substance inside a container often are references to the amount of that substance inside the container, e.g., "*the steam in the boiler*". However, we cannot allow all containment statements to be quantities, e.g. "*I am in a state of shock*" is not a quantity statement. We distinguish between these cases by requiring the entity to be an instance of `ChemicalSubstanceType`. There are yet more complex cases, even for quantities. Some quantities are implied, e.g., "*the hot brick.*" Adjectives like *hot* often modify a specific quantity type, so such cases are handled by looking for quantity slots (e.g. `temperatureOf`) and connections between values (e.g. "*hot*") and quantity types (e.g. `Temperature`).

```
(<== (introducesQPQuantityFrame
       (PresentationEventFn ?sid ?nevent)
       ?qframeid)
   (quantityTypeOfEntityFound ?sid
     ?qtype ?quantity ?entity ?etype)
   (buildsQPFrame ?sid ?qtype ?quantity
    ?entity ?etype ?qframeid ?nevent))

(<== (quantityTypeOfEntityFound
       ?sid ?qtype ?qres ?eres ?etype)
   (contextIndependentDrsFor ?sdrs ?sid)
   (getAllPotentialBinPreds ?sid ?entity
     MeasurableQuantitySlot ?pred
     WorldLikeOursCollectorMt)
   (ist-Information (DrsCaseFn ?sdrs)
     (?pred ?entity ?quantity))
   (isaIn2ndOrderCollectionOf
       ?qtype QuantityType)
   (ist-Information (DrsCaseFn ?sdrs)
      (isa ?entity ?etype))
   (resolvedVar ?sid ?entity ?eres)
   (resolvedVar ?sid ?quantity ?qres))
```

Table 2: Example of rules involved in computing narrative function.

Table 2 illustrates some of the rules for introducing quantities. The first rule constructs the presentation event and frame (via bindings returned and the rules which implement `buildsQPFrame`, respectively). The seven rules for `quantityTypeOfEntityFound` (we only show one due to space limitations) use EA NLU's syntactic and semantic analysis to identify such frames. The first antecedent binds context variables for the current sentence and its default discourse representation structure. The second antecedent finds all bindings for entities which participate in statements whose predicate is an instance of the concept `MeasureableQuantitySlot`, for which there are over 300 relationships in the knowledge base, such as `temperatureOfObject`, relative to the logical environment representing common background knowledge (`WorldLikeOursCollectionMt`). Abductive inference happens with the `ist-Information` antecedents, which cause the reasoning system to explore collections of assumptions that would generate potential bindings. The relationship `isaIn2ndOrderCollection` ensures that there is a concept which is an instance of the (higher-order) concept `QuantityType`, such as `Temperature`, of which the value of the slot (?quantity) is an instance. Finally, the `resolvedVar` antecedents invoke coreference resolution, returning either the prior referent or the current discourse variable as the binding for their second argument.

A hallmark of natural language is that it often provides only partial information about a situation, which is why frame representations are so useful in semantics. Even though higher-order frames are sought after lower-order frames, incremental processing means that we must be able to merge information across sentences. Consider the following two sentences which, together, entail a quantity transfer: *"Heat flows from the hot brick. Heat flows to the cool ground."* Understanding this discourse fragment requires recognizing that the *flow* event in both sentences is the same, which also suggests that the *heat* is the same, after which the implied direct influences can be recognized. Kuehne (2004) used antecedent rules to merge quantity frames both within and across sentences. Instead, we extended the abductive coreference algorithm of Tomai & Forbus (2009) to include verb coreference, by searching for multiple verbs that have the same event type and root.

An analysis of a broader range of texts revealed an interesting assumption implicit in Kuehne's analysis of direct influences. The sentences above would have resulted in a single rate parameter, i.e. the rate of transfer of heat from the brick to the ground is the same. However, consider the following sentences: *"Heat flows from the hot coffee. The heat flows to the cold ice cubes and the cool mug."* Here the *flow* events may be coreferents, but assuming energy conservation, the rate of heat transfer from the coffee cannot be the same as the rate of transfer to the ice cubes and to the rate of transfer to the mug. Thus we do not merge coreferent rates: Another direct influence could always come along in the next sentence. Instead, we postpone such closed-world assumptions to subsequent processing.

Evaluation

The system performs accurately on all 8 examples from Kuehne (2004). We further evaluated system performance on the first nine simplified paragraphs from chapter two of a science book intended for general readers (Buckley, 1979). The sentences were taken from the same corpus previously used by Barbella & Forbus (2011) and follow their simplification paradigm. That is, syntax is simplified by breaking complex sentences into multiple sentences (roughly the level of the grammar found in middle-school reading comprehension books), but leaving the vocabulary intact whenever possible. The corpus was hand annotated for QP frames, and the system compared to this gold standard. Only frame types used by Kuehne (2004) were evaluated. Thus, we did not include topology frames, generic type-level frames, or limit points. After each paragraph, the reference context was cleared.

Of the 144 tagged frames in the corpus, our system correctly constructed 65. There were 23 extraneous, partial, or incorrect frames generated as well. This gives us a precision of .74 and a recall of .45. The F1 harmonic mean was .56. We view this as a respectable start, and given that this is a new problem, a good baseline against which to judge future efforts.

Our analysis suggests that there are two sources of errors. First, errors in entity coreference resolution leads to duplicate low-level quantity frames and incomplete higher-order influence frames. A second source of errors was modal modified exchange verbs, as in: *"You can buy heat from a gas company"*. Given that exchange relationships involve two distinct quantity transfer relationships, failures on these sentences significantly reduce performance.

Narrative Functions for Type-level Influences

Recently QP theory was expanded to include *type-level influences* (Hinrichs & Forbus, 2012). Type level influences are a form of higher-order qualitative reasoning, expressed in terms of causal relationships between predicates and concepts, rather than specific individuals. Type-level influences can provide significant benefit in large-scale domains and planning tasks. For example, the strategy game Freeciv[3], an open-source version of the classic computer game Civilization, provides a rich environment for experimenting with how qualitative reasoning can be used for modeling the kinds of reasoning and learning involved in understanding economics, strategies, and tactics. In Freeciv players build civilizations by founding cities, researching new technologies, improving the land around their cities, and building settlers to found new cities, to expand their civilization further. Such games are far more complex than chess, for example, and require many hours to learn.

[3] http://freeciv.wikia.com/wiki/Main_Page

Interestingly, important advice can often be expressed in language whose semantics is well captured by type-level influences. For example, the statement

"Adding a university in a city increases its science output"

can be formally expressed via this type-level influence:

```
(positivelyDependsOn-TypeType
  (MeasurableQuantityFn cityScienceTotal)
  FreeCiv-City FC-Building-Univerity
  cityHasImprovement)
```

That is, the science output of a city (which is a measurable quantity, i.e. one that can be read out of the simulator) can be positively affected by adding an improvement to the city (by achieving a `cityHasImprovement` statement) where the type of improvement is a university (`FC-Building-University`).

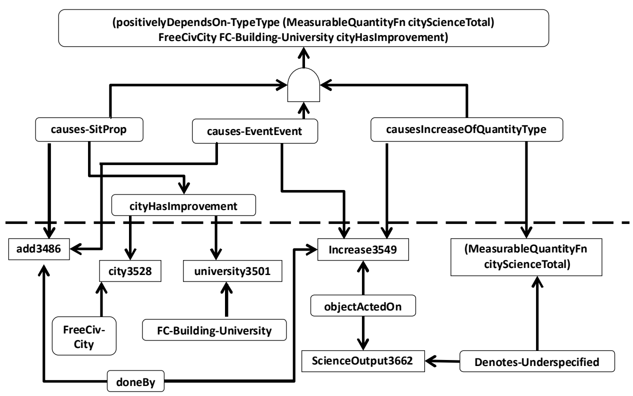

Figure 2: Type-level inference derivation from language analysis

To extend narrative functions to handle such type-level influences, we added one new type of narrative function, `IntroductionOfFCRelation`, indicating that new game-relevant information was detected. The new detection rules were of two types. The first extracts a layer of causal relationships from the events found in the linguistic analysis. For example, the sentence above includes two events, one referred to by "adding" and the other referred to by "increases". Since there is a `doneBy` relationship produced by the parser that links the two events, the narrative function rules infer a causal relationship between them. That is, the `Incorporation-Physical` event causes the `IncreaseEvent` event. The second type of detection rule looks for causal patterns that suggest an influence at work. For example, if an event causes some statement to be true, and the same event is the causal antecedent of a quantity change event, then that suggests that statement is the condition to use in the type-level influence.

In addition to new narrative function rules, additional facts were added that biased the scoring system for abduction to prefer solutions containing type-level influences and narrative functions. For example, the interpretation of "adding" above to mean the arithmetic operation applied over two numbers did not give rise to causal connections that allowed an influence to be produced, leading the system to automatically prefer physical incorporation as the intended meaning of the word.

Figure 2 depicts a partial dependency structure showing how the influence above was inferred from the analysis of the sentence. The entities and relationships below the dashed line were produced by the parser, while the statements below it were produced by the narrative function rules. Notice that upper layer consists of very general causal relationships. We suspect that this structure will be very general: The variations in the specifics of language might be handled by rules that produce these general causal relationships, while the more complex narrative functions can be captured by patterns that are truly domain-independent. Whether or not this scales is, of course, an empirical question.

When viewed as advice, is this type of information useful? To find out, we ran a Companion (Forbus et al 2009) with and without the following pieces of advice:

- Adding a granary in a city increases its growth rate.
- Adding a research lab in a city increases its science output.
- Adding a library in a city increases its science output.
- Adding a university in a city increases its science output.
- Irrigating a place increases food production.
- Mining a place increases its shield production.

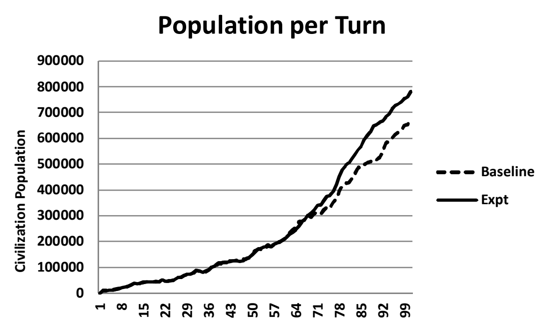

Figure 3: Population growth improves with advice

We created different maps by saving the game on the initial turn, with the default map settings. A Companion then played for 100 turns on each map under two different conditions. In the baseline condition, the Companion used a qualitative model that was previously learned, entirely through demonstration. In the experimental condition, the Companion also had access to the type-level qualitative influences obtained by reading the six sentences above. Figures 3 and 4 show the difference in the two conditions, averaged over 10 games. The improvement in population growth (Figure 3) is due to the effect of irrigation, while the improvement in science output (Figure 4) is due to the other improvements. The improvement in science output is statistically significant ($p < 0.041$), while the improvement in food production is not quite significant ($p < 0.06$). We believe that the full effect of irrigation on population takes longer than 100 turns to manifest, and may also become significant over longer portions of the game. Regardless, the type-level influences extracted from language altered the behavior of the system such that it tasked worker units with creating irrigation and built libraries in its cities. This is encouraging evidence for the utility of type-level influences, expressed via natural language, for giving advice to cognitive systems.

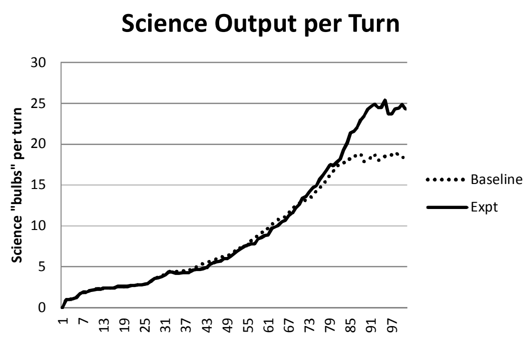

Figure 4: Science output improves with advice

Related Work

While our computational approach to constructing narrative function is not intended as a process-level cognitive simulation of reading, it is compatible with the results of (Graesser et al., 1994) in that inferences concerning the interconnections implied by the text are constructed on-line, during the understanding process. The representation of narrative functions developed in Tomai (2009) goes beyond work on story grammars (e.g. Trabasso et al., 1984) in that it supports multiple points of view, and Tomai's abductive algorithm supports both interleaving of patterns and allowing story elements to participate in multiple narrative functions, which are difficult to handle in story grammars.

Our abductive approach differs from Hobbs (2004) by only allowing assumptions about choice sets, rather than arbitrary domain assumptions. Furthermore, Hobbs (2004) favored minimal assumptions in their cost function while our system relies more heavily on type-level reasoning to select among competing choices. It also differs from Ovchinnikova (2012), which uses a knowledge base extracted from WordNet and FrameNet and uses lexical knowledge to weight abductive inferences. In contrast, our approach focuses on how discourse and narrative goals can guide abductive inference from the top down. That said, the two approaches are not mutually exclusive and exploring their combination might be worthwhile. Finally, Blythe *et al* (2011) investigated an implementation of weighted abduction using Markov logic networks, but was not a top-down narrative algorithm, like ours.

The closest work in using natural language to improve strategy game performance is that of Branavan et al (2011), who used dependency parses of the Civilization 2 manual to suggest linkages between game concepts to bias a Monte Carlo learner. They limited their experiments to a game board $1/4^{th}$ the normal size, which facilitated experiments, but also simplified the problem, since many more challenging aspects of the game were factored out. We are tackling the full complexity of the game: At 100 turns a real game is just starting, whereas theirs were over. Our approach requires less text, i.e. just six sentences, completely understood, leads to significant performance gains. Moreover, our player uses qualitative reasoning and achieves immediate learning improvements, whereas their system used the game engine to do massive (8-way parallel lookahead) computation. Most complex dynamic systems people deal with do not have accurate simulations available, and thus we believe our approach will scale better to more real-world applications.

Conclusions & Future Work

We have shown evidence that the concept of narrative function can be used to understand texts whose meaning include information expressible via QP theory. It performs as well on the original examples of Kuehne (2004), but also does respectably well on material from a science text. Moreover, we have shown that this approach can be used to learn advice from language whose meaning can be captured via type-level influences.

While much improved over Kuehne (2004), the biggest limitation of the current system remains coverage. We plan to address this by several lines of future work. First, we plan to expand the coverage of instance-level qualitative descriptions, to handle the range of QP-bearing language found in science books. Second, we plan to expand the coverage of type-level qualitative descriptions, to handle the sorts of descriptions of continuous processes, quantities, and relationships found in both science books and in discussions of planning and strategies involving dynamical systems (for which Freeciv is a useful laboratory). Third, we plan to expand the coverage of narrative functions to handle the rest of the material in such texts. Introducing new principles, problem-solving strategies, and examples, for instance, are common types of narrative functions in such texts. Fourth, our current abduction system does not support backtracking well, nor does it gracefully incorporate evidential reasoning or the use of analogical abduction. We are currently designing a new abduction system aimed at overcoming these limitations.

Acknowledgements

This research was sponsored by the Intelligent Systems Program of the Office of Naval Research (Grant N00014-07-1-0919) and the Air Force Office of Scientific Research (Award FA2386-10-1-4128).

References

Allen, J. F. (1994). Natural language understanding (2nd Ed.) Redwood City, CA: Benjamin/Cummings.

Branavan, S., Silver, D., & Barzilay, R. (2011). Non-Linear Monte-Carlo Search in Civilization II. *Proceedings of IJCAI 2011.*

Buckley, S. (1979). *Sun Up to Sun Down: Understanding Solar Energy*. McGraw-Hill.

Barbella, D. and Forbus, K. (2011). Analogical Dialogue Acts: Supporting Learning by Reading Analogies in Instructional Texts. Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011), San Francisco, CA.

Barthes, R. (1977). Introduction to the structural analysis of narratives (S. Heath, Trans.). In *Image, music, text.* New York: Hill and Wang.

Blythe, J., Hobbs, J.R.,Domingos, P. Rohit, K.J. & Mooney, R.J. (2011).Implementing weighted abduction in Markov logic. In Proceedings of the Ninth International Conference on Computational Semantics (IWCS '11). Association for Computational Linguistics, Stroudsburg, PA, USA, 55-64.

Fillmore, C. J., Wooters, C., & Baker, C. F. (2001). Building a Large Lexical Databank Which Provides Deep Semantics. In Proceedings: Pacific Asian Conference on Language, Information, and Computation, Hong Kong, China.

Forbus, K. D. (1984). Qualitative Process Theory. Artificial Intelligence, 24, 85-168.

Forbus, K., Klenk, M., and Hinrichs, T. 2009. Companion Cognitive Systems: Design Goals and Lessons Learned So Far. *IEEE Intelligent Systems*, 24(4), 36-46.

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

Grishman, R., Macleod, C., & Wolff, S. (1993). *The COMLEX Syntax Project*. Ft. Belvoir: Defense Technical Information Center.

Hinrichs, T. and Forbus, K. 2012. Learning Qualitative Models by Demonstration. *Proceedings of QR2012.*

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

Kamp, H., & Reyle, U. (1993). *From discourse to logic: Introduction to model-theoretic semantics of natural language.* Boston, MA: Kluwer Academic.

Kuehne, S. E. (2004). Understanding natural language descriptions of physical phenomena. Doctoral dissertation, Northwestern University, Evanston, Illinois.

Labov, W., & Waletzky, J. (1966). Narrative analysis: Oral versions of personal experience. In J. Helms, (Ed.), *Essays on the Verbal and Visual Arts* (pp. 12-44). Seattle: University of Washington Press.

Ovchinnikova, E. (2012). Integration of World Knowledge for Natural Language Understanding. Atlantis Press, Springer.

Tomai, E. (2009). A Pragmatic Approach to Computational Narrative Understanding. Doctoral dissertation, Northwestern University, Department of Electrical Engineering and Computer Science, Evanston, Illinois.

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

Trabasso, T., Secco, T., & van den Broek, P.W. (1984). Causal cohesion and story coherence. In H. Mandl, N.S. Stein & T. Trabasso (Eds.), *Learning and comprehension of text* (pp. 83-111). Hillsdale, NJ: Lawrence Erlbaum Associates.