

Using Narrative Function to Extract Qualitative Information from Natural Language Texts: A Preliminary Report

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 (Kuehne 2004) showed that frame-based representations of qualitative process theory constructs could indeed be extracted from natural language texts. Kuehne’s approach relied on the parser recognizing specific syntactic constructions, which has 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 type-level qualitative representations (Hinrichs & Forbus, 2012) can be extracted from text, and used to improve performance in playing a strategy game.

1 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 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 *frame elements*, i.e. slots in the frame. For example, the notion of qualitative proportionality is captured by an Indirect Influence frame, which includes the following frame elements:

- *Constrainer*: The antecedent quantity of the causal relationship

- *Constrained*: The quantity being constrained by this relationship
- *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. Here is an example of such a pattern:

AS <Quantity1> <Change1>, <Quantity2>
<Change2>.

“As the **temperature** of the steam **rises**, the **pressure** inside the boiler **rises**”

In the example above, the constrainer would be a quantity frame representing the temperature of the steam, the constrained would be a quantity frame representing the pressure inside the boiler, and since both of the changes are a form of increasing, the sign would be 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. While this was a successful proof of concept demonstration, when trying to scale this up for use in systems that learn by reading, we discovered several limitations. First, the use of syntactic patterns significantly limited coverage. Second, the antecedent rules used to merge coreferential frames did not scale well to larger texts.

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 narrative functions, and that this approach already captures almost all of the range of examples handled previously. 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 future work.

2 Narrative Function and Abduction

When people read, they try 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 provides a 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 (2009) showed that narrative functions could be used in understanding natural language texts such as fables and the materials used 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. Section 3 describes these new narrative functions. But first, we provide some relevant background about the natural language system, EA NLU.

2.1 EA NLU and Choice Sets

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) and ResearchCyc¹ for syntactic processing. It uses representations from ResearchCyc for its semantics, including 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:

Q: "How many children does Mary have?"

A: "She has 3 kids."

The term *kid* is ambiguous. It could be a child, or it could be a baby goat, as these choices from the KB indicate:

```
(isa kid5283 HumanChild)
(isa kid5283 (JuvenileFn Goat))
```

Here *kid5283* is a *discourse variable*, an arbitrary individual introduced to represent whatever it is that "kid" refers to. This is an example of a *word sense* choice set. The other kind of choice set produced by EA NLU concerns parsing choices, e.g. where a prepositional phrase should be attached. Semantic interpretation involves selecting an element from each choice set implied by the linguistic analysis of the sentence. This can be quite complex: For example, choices in some choice sets might imply the existence of further choice sets to be considered. In general, semantic interpretation is an unsolved problem. Strategies like backtracking search have been tried, but they flounder on the large number of possible interpretations. Interestingly, psycholinguistic research suggests that people are quite rapid readers, and seem to do long-range backtracking very rarely. There are many possible explanations for this, including performing evidential reasoning to select the most likely choices. Another source of constraint is context, which provides expectations. Here, the first interpretation of "kid" would be more sensible, because "She" presumably is co-referential with "Mary", and since the question spoke of "children", we might assume that Mary is human. Therefore (trans-genetic experiments notwithstanding) Mary's child is most likely human. This choice supports the second statement being an answer to the first, which is an example of narrative function in action.

2.2 Abduction

EA NLU uses a novel query-driven abduction process to provide top-down guidance to the process of semantic interpretation. Abduction is inference to a plausible explanation. That is, if $P \Rightarrow Q$, then an explanation for Q being observed is that P is true. Obviously there could be other explanations for Q , so abduction is not deductively valid, and relies on heuristics for estimating the plausibility of abductive assumptions. Abduction has long been used in semantic interpretation (Hobbs 2004), but it tends to be intractable as the number of statements grows. Tomai (2009) showed that by using top-down expectations, e.g. looking for a moral choice, the complexity of abduction over a discourse could be greatly reduced, since many potential choices could simply be ignored.

The abduction mechanism in EA NLU only makes assumptions about what choices should be made from the choice sets presented by linguistic analysis. It is driven by queries, which are generated based on overall context of the task as well as specifics in the data. To identify the set of queries to be made, it first does a query of the form

```
(queryForInterpretation ?o ?q)
```

?q is a query that should be made in the interpretation context for the current sentence. ?o is an integer that provides advice about the ordering of queries. All queries with lower values for order will be done before any query

¹ <http://www.cyc.com/platform/researchcyc>

with a higher value for order. Thus, for example, the rules searching for influences can be assured that any quantity information already existing in the discourse will have been found. We call this mechanism *query for questions*.

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. This approach differs from that of more lexically oriented abductive NLU systems such as (Ovchinnikova 2012). Ovchinnikova’s abductive NLU system operates over a knowledge base extracted from WordNet and FrameNet and uses lexical knowledge to weight abductive inferences. Our approach instead focuses on how discourse and narrative goals can guide abductive inference. Of course, the two approaches are not mutually exclusive and future work will certainly focus on incorporating more word and sentence level pragmatic knowledge.

2.3 Representing Narrative Functions

The connection between a piece of sentence content and its role in the narrative is expressed via

```
(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 substituted into each query processed by the query for questions mechanism outlined above, 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 and the types are from an ontology outlined below.

3 Finding QP Frames via Abduction of Narrative Function

This section outlines the narrative functions for QP frames that we have developed, and summarizes some important properties of the rules that derive them from the natural language analysis of texts.

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

3.1 QP Specific Narrative Functions

For each QP Frame type, we introduce a category of narrative function (see Table 1). For example, the following query kicks off the search for quantities:

```
(queryForInterpretation 0
 (narrativeFunction
  (PresentationEventFn :REPLACE-SID ?event-id)
  ?quantity-frame
  IntroductionOfQuantityEvent))
```

The order information associated with a query is used to organize the computation so that higher-order narrative functions are only sought after their potential constituents have been identified. For example, participants and consequences of continuous processes are sought after quantities have been found, and also after ordinal relationships have been detected, e.g.

```
(queryForInterpretation 3
 (narrativeFunction
  (PresentationEventFn :REPLACE-SID ?event-id)
  ?process-frame-role IntroducesProcessRole))
```

In addition to the QP Frame types proposed by Kuehne (2004) we created a frame for describing topological constraints on a system such as connections, interruptions, and paths. 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. An example would be:

“Cylinder A1 is connected to Cylinder A2 by a pipe.”
“Water flows from Cylinder A1 to A2.”

3.2 Basic QP Frame Extraction

Solutions to narrativeFunction are found via Horn clause rules which are similar to Prolog rules. They are different in that there is no notion of cut and all solutions are found. 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 a statement with the predicate `in-UnderspecifiedContainer`. 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, 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 QType is the continuous quantity type.

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`).

3.3 Discourse Level QP Frames

Process Frames and Quantity Transfer Frames both require information from lower-level QP frames such as quantities. Thus narrative functions for these frames are sought after low-level queries have been completed. However, a hallmark of natural language is that it often provides only partial information about a situation. Thus not all of the constituents may be available, which is why frame representations are so useful in semantics. For example, we may know that there is a process going on based on the use of a process verb, but the sentence may not provide enough information to generate direct influences or qualitative proportionalities.

Another complexity is that higher level frames often combine 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 the quantity transfer frame implicit in the above sentences requires recognizing that the *flow* event in both sentences is the same. This would also suggest that the *heat* is the same. Only then do the two direct influences implied by the pair form source and destination assertions.

Kuehne (2004) used antecedent rules to merge quantity frames both within and across sentences. Instead, we extended the abductive coreference algorithm of (Tomai 2009) to include verb coreference. This works 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.”

In the above, the *flow* events may be coreferent.

However, 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. Because of this, while we merge coreferent events, we do not merge coreferent rates: Another direct influence could always come along in the next sentence. Instead, we assume that downstream reasoning should be used to introduce such assumptions, based on closed-world assumptions over the material being read.

3.4 Evaluation

The system was evaluated using eight gold-standard QP examples from (Kuehne 2004). The texts covered all possible types of QP frames and several were multiple sentences long. The QP frames produced by the two systems were compared. For example, Figure 1 is a graphical depiction of the QP frames produced for the sentence “Heat flows from the brick.”




Figure 1: QP Frames for “Heat flows from the brick”

Currently, the system performs accurately on seven of eight examples. The incorrect example fails due to errors in coreference resolution. The other limitation is that we do not currently implement the Preconditions frame element for process frames. Other than those two differences, the results are compatible with Kuehne (2004).

4 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 Freeciv², 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. 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-University
 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 which is a University (i.e. achieving a cityHasImprovement statement relating a city in a Freeciv game with an instance of the concept of university in Freeciv).

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 statements were made 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 in blue were produced by the parser, while the statements in yellow were produced by the narrative function rules. Notice that the yellow 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.






Figure 2: Type-level inference derivation from language analysis

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.

² http://freeciv.wikia.com/wiki/Main_Page

- Adding a university in a city increases its science output.
- Irrigating a place increases food production.
- Mining a place increases its shield production.



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. This is encouraging evidence for the utility of type-level influences, expressed via natural language, as a means of giving advice to cognitive systems.

5 Conclusions and 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 almost as well on the original examples of Kuehne (2004), but also can be used to learn advice from language whose meaning can be captured via type-level influences. However, we view these results as preliminary because of limited coverage to date.

We plan to explore several directions in future research, most of them concerned with expanding different aspects of coverage. First, we need to expand the coverage of instance-level qualitative descriptions significantly, to

handle the range of QP-bearing language found in science books. Second, we need to expand the coverage of type-level qualitative descriptions, to handle the 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 need 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 is limited, in that it 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 that we hope will overcome these limitations.

6 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).

7 Bibliography

- Allen, J. F. (1994). *Natural language understanding* (2nd Ed.) Redwood City, CA: Benjamin/Cummings.
- 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.
- Grishman, R., Macleod, C., & Wolff, S. (1993). *The COMLEX Syntax Project*. Ft. Belvoir: DTIC.
- 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.
- 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.