Unifying Instance-level and Type-level QP Frames for Natural Language Understanding

Kenneth D. Forbus and Thomas Hinrichs

Abstract. One important role for qualitative representations is as a constituent of natural language semantics. The incremental nature of language means that information about situations and models typically arrives piece by piece, and must be assembled in order to create a qualitative model that can be reasoned with. Our prior research created distinct QP frame systems for instance-level and type-level qualitative models. This is inelegant as well as problematic, since the choice of constructing instance versus type level models ought to be made based on the text and task context, not a priori. This paper describes the design of a unified QP frame system, where most of the frame contents are agnostic with respect to whether the final description will be instance-level or type level. Clues from the NLU system's analysis plus context will determine whether the final model constructed is instance level, type level, or a mixture of the two.

1 Introduction

Language is a primary means of human cultural transmission, providing concepts that help organize our mental models of the world. Since qualitative representations are aimed at expressing human mental models of continuous domains, it seems natural that qualitative representations should be part of computational models of natural language semantics. Qualitative process theory [6] has been used to construct qualitative representations of specific scenarios from natural language paragraphs [15] and also to learn qualitative models by reading (e.g. [18][20]). Our approach is to understand how to perform high-precision understanding of natural language, in order to create cognitive systems that can construct explicit qualitative models from reading and dialogue. This is no longer the only approach being explored. As researchers in machine learning and computational linguistics begin to explore how to learn commonsense knowledge from text, they are starting to tackle overlapping problems. In some cases they explicitly use ideas from qualitative reasoning, e.g. the AI2 QuaRel dataset [22] concerns comparative analysis questions and a set of qualitative proportionalities are supplied as part of the background knowledge. Other AI2 datasets[2] (e.g. ProPara, ROPES, SciTail, ARC) involve phenomena explored in qualitative reasoning research, but with multiple other forms of knowledge involved as well. Most of these approaches use distributed representations without any attempt to construct conceptual representations. This can lead to surprisingly good performance on multiple choice questions for which the system was trained. However, there is evidence that such systems are actually quite brittle, susceptible to noise and their performance outside the original training set is hard to predict [12][16]. Moreover, they are not able to handle questions requiring answer generation, as opposed to multiple choices. They do not provide explanations for their results, nor can they take advice and correction when they make mistakes. To provide these capabilities, we believe that a high-precision understanding process which produces explicit internal representations (including qualitative representations) is needed. This is what we are exploring through our Companion cognitive architecture [8].

This paper describes how we are extending our representations that integrate qualitative representations into natural language semantics to handle a broader range of phenomena, such as the AI2 datasets, but also learning by reading books. We begin by summarizing how we have used different frame representations to provide an intermediate level of representation between natural language and qualitative models, and why a new, unified QP frame representation is needed. The key properties of such a representation are discussed in Section 3, including the frame types and slots required. Section 4 illustrates using examples, Section 5 reports on an implementation in progress, and Section 6 wraps up and outlines future work.

2 QP Theory and NL Semantics

Qualitative representations have, from their beginnings, been inspired by distinctions found in natural language. The utility of concepts such as signs of derivatives, for instance, are motivated in part by the common terms "increasing", "decreasing", and "steady." Kuehne's work [15] identified a set of language patterns that could be mapped directly to constructs of QP theory. For example, the comparative correlative construction (e.g. "the bigger they are, the harder they fall") can be translated into a qualitative proportionality, as can phrases like "depends on" (e.g. "level depends on volume"). This includes information about continuous processes, e.g. "Heat flows from the hot brick to the cool ground." Language often provides information in pieces that must be assembled to construct an assertion-based qualitative model. For example, "The brick is hot and the ground is cool. This causes the heat to flow from the brick to the ground." The condition of the process, an ordinal relationship involving temperature, must be combined ("this causes") with the occurrence of the heat flow process. Such incrementality is a hallmark of language. This is why most accounts of natural language semantics use some form of *frame representation* [5][21], where some parts of the text introduce frames that are filled in by the meanings from other parts of the text.

[1] Qualitative Reasoning Group, Northwestern University, 2133 Tech Drive, Evanston, IL, 60201, USA. forbus@northwestern.edu

[2] See https://leaderboard.allenai.org/ for details.

Most work in qualitative reasoning involves logically quantified model fragments which are instantiated to create models of specific scenarios (e.g. [3]). Kuehne's original frame system constructed instance-level representations, i.e. descriptions about a specific scenario. The difference between a specific scenario and general principles as expressed in language can be complicated, as the difficulty in human judgments concerning whether or not a statement is generic indicate. Abstract examples are often used in explanatory texts, for example, where a reader either abstracts out a logically quantified description, or stores an instance-level description which can then be applied via analogy to new situations (e.g. [7]). Exploring qualitative reasoning in constructive dynamic domains, such as strategy games, where large numbers of entities exist, and the set of entities changes rapidly (e.g. expansion of a civilization, warfare), has led to the formulation of type-level qualitative models [11], which provide concise descriptions to support larger-scale qualitative reasoning. As we explored using natural language advice [18] and learning qualitative models by reading [20], we built up an alternate set of QP Frames around type-level qualitative representations. While these type-level frames have proven to be useful, there are still circumstances where instance level representations are more appropriate. Currently we, the experimenters, choose which system is to be used, which is inelegant. There is a large degree of commonality between these two systems, and so it makes sense to integrate them.

How should a system decide which kind of qualitative model to construct from language, either when reading or participating in a dialogue? What are the sources of ambiguity with regard to the choice of instance-level versus type-level modeling? Both involve the same kinds of causal laws and conceptual packaging of these laws and their applicability conditions. What varies is the nature of the entities over which they are applied. Consider this statement about a strategy game's dynamics: "Production depends on the size of the population." Prior work on QP/language mappings suggest that this is a qualitative proportionality, and the quantity types involved are production and population. But in such games, population is both a property of cities and of an entire civilization. Which usage of population is intended here? The sentence itself is not telling us. Such sentences can be constructed, e.g. "Production in a city depends on the size of the city's population.". However, Gricean postulates on communication suggest that what is known in context should not be belabored [9], hence seemingly obvious information should left out. Another issue concerns which kind of qualitative model is being described. This is a type-level, general interpretation of the sentence. Could it have an instance-level interpretation as well? Yes, since context also matters. If the original sentence about production were a response to a question, e.g. "What does the production in Boston depend upon?", then this sentence is more naturally viewed as describing an instance-level statement about causality in the situation being asked about, i.e. Boston. Thus higher levels of language understanding need to be invoked to ascertain whether a type-level or instance-level qualitative model needs to be formulated, given the context. That is why we need a single system of QP frames which defers this decision to a later phase which can take more context into account when formulating assertion-based qualitative models for reasoning.

3 A unified QP Frame representation

We use as our starting point the OpenCyc ontology [16] and the FrameNet frame semantics representations [5], which we have linked in our NextKB knowledge base[3]. NextKB also includes a large vocabulary English lexicon, and support for qualitative reasoning, textbook problem solving, and visual reasoning. The mappings between QP theory constructs and language that we build upon here are described in [19]. The Companion natural language understanding system uses narrative functions to construct higher-level interpretations of language [23]. It uses Discourse Representation Theory [12] to handle contexts, counterfactuals, and logical and numerical quantification.

In our new intermediate QP frame representation, the difference between instance level representations and type-level representations hinges on the kinds of entities involved. In QP theory, entities appear in two distinct roles. First, entities are the things whose continuous properties are the quantities constrained by causal laws and whose changes over time are reasoned about. Second, entities are participants in model fragments, whose bindings then provide them as arguments to preconditions and consequences. As long as relations taking entities as arguments (e.g. slots in a frame) can handle both individual objects (e.g. Boston, in a strategy game) and types (e.g. City, again in a game), the rest of the frame elements can be the same. We believe that this will greatly simplify the extraction of QP frames from text, while making the construction of an assertional representation for reasoning somewhat more complicated.

This section describes the current design for our new unified QP frame representation. Following Hayes [10] and the Cyc project, frames are represented as sets of predicate calculus assertions. That is, the following are equivalent[4]:

(isa *<frame token>* *<type of frame>*)
(*<slot1>* *<frame token>* *<value1>*)
(*<slot2>* *<frame token>* *<value2>*)

<frame token>
Type: *<type of frame>*
<slot1>: *<value1>*
<slot2>: *<value2>*

The second notation is more traditional in the literature on frames, and is more compact on the page. The assertion-based implementation means that systems can freely intermingle assertional statements and frames, as needed.

We proceed via examining, for each construct of QP theory, what the corresponding frame representations are and what linguistic clues are available to provide evidence for type-level versus instance-level representations.

[3] http://www.qrg.northwestern.edu/nextkb/index.html is available under a Creative Commons Attribution-only license.
[4] Frame systems often include specifications of memory retrieval, i.e. all of the slots are retrieved when the frame is, a consideration not captured by Hayes' logical encoding.

3.1 Quantity Frames

Quantity frames represent information about a quantity. The slots include

- `quantityType`: The kind of quantity involved, e.g. pressure.
- `quantityEntity`: An object or type, e.g. Boston, City.
- `valueOf`: A numerical or symbolic value, e.g. `(HighAmountFn Production)`[5]. Units are encoded as non-atomic terms, e.g. (meters 3).
- `dsValue`: The sign of its derivative, either `-1`, `0`, `1`, or unknown.

Only `quantityType` is required. That is, a phrase like "The production depends on…" should not require choosing an entity in order to construct a QP frame to capture what is being expressed in that phrase. To fully interpret the phrase, the entity must be determined, of course, but that can be taken care of by a later process that assembles models from frames. `valueOf` and `dsValue` can have multiple values, e.g. when both a numerical and symbolic value are known (e.g. "…300 degrees C, which is really high."), or partial information about change is known (e.g. "is non-decreasing").

3.2 Ordinal Frames

Ordinal frames represent information about ordinal relationships. The slots, all single-valued, include:

- `quantity1`: A quantity frame
- `quantity2`: A quantity frame
- `ordinalReln`: An ordinal relationship, one of `lessThan`, `greaterThan`, `equalTo`, `lessThanOrEqualTo`, etc.

At least one quantity and the ordinal relation must be specified, e.g. "My package is heavier." implies another package, but that presumably will be added by combining meanings across sentences.

3.3 Influence Frames

Influence frames have the following single-valued slots:

- `constrained`: A quantity frame
- `constrainer`: A quantity frame
- `influenceType`: `Direct` or `Qprop`
- `influenceSign`: + or −
- `sourceOf`: A model fragment frame.

The patterns in [15][19] indicate whether `Direct` or `Qprop` should hold. This differs from prior schemes in having a single frame type for both kinds of influences, rather than having a subclass of frames for each kind. This is because some patterns, e.g. "X affects Y" actually are ambiguous with respect to direct versus indirect (qprop) influences, and hence `influenceType` is optional. On the other hand, we have not yet found examples where just one quantity naturally occurs – asking "What constrains the temperature?" implies that bindings for the constrainer are sought.

[5] The OpenCyc ontology has an elegant system of symbolic qualitative values, as well as units for numerical values.

3.4 Model Fragments and Continuous Processes

Model fragment frames have the following slots:

- `participantOf`: Multiple values, each an entity or type.
- `conditionOf`: Multiple values, statements and/or ordinal frames that must hold for the model fragment to be active.
- `consequenceOf`: Multiple values, statements and/or that hold when the model fragment is active
- `statusOf`: active or inactive, depending on the conditions

Continuous process frames are a subclass of model fragment frames, with an additional slot

- `directInfluenceOf`: Mulitple values, all direct influence frames. Direct influence frames cannot appear as arguments to `consequenceOf` slots, to enforce the sole mechanism assumption of QP theory [6].

3.5 Other Frame-like Relationships

In linguistics, events and configurations are often represented by frames, again due to the need to incrementally accumulate information when understanding language [5]. The NextKB representations for events are built upon OpenCyc, which uses the same method for representing frames. Moreover, we have also found frames useful for configuration information, e.g. paths, which in most QR work are represented by a single assertion. While sometimes an entire path is encoded in a single sentence (e.g. "Heat flows from the surface of the hot brick to the cold ground."), configurations are often built up incrementally in language as well. The common abstractions of paths and containers are two examples of such structural concepts.

4 Examples

To get a sense of how these representations will be used, we work through two simple examples. The first is the production example from Freeciv. The sentence "Production depends on population." would be encoded as a single influence frame, with two quantity frames as arguments:

```
Q1:
  quantityType: Production
Q2:
  quantityType: Population
Inf1:
  constrained: Q1
  constrainer: Q2
  influenceType: Qprop
  influenceSign: +
```

Coreference resolution occurs across entities in a discourse. If the previous sentence had been "You should increase production in Boston." the qualitative component of that sentence would be:

```
Q0:
  quantityType: Production
  quantityEntity: Boston
```

```
(dqValue Q0
  (BeforeAfterMappingFn increase3) 1)
```

Where dqValue indicates a differential qualitative value [24], comparing across the before and after states of the increase event referred to in the sentence (represented by the discourse variable increase3), with "increase" mapping to 1. (The use of analogical mappings to perform cross-situation alignment was first explored in [14].) We leave out the semantics of "You should" since it does not impact the QP frame representation. Coreference resolution will merge Q0 and Q1. Should quantityEntity information automatically be propagated to Q2? In this case that would be correct, but it needn't be. Consider "Production depends on tax rate." Tax rates are set with respect to the entire civilization, not the specific city, and that propagation would thus be incorrect. The system must puzzle this out, based on what it knows, or ask the user for clarification. For example, there are only two types of population quantities in the dynamics model, the population of a civilization and the population of a city. The former is qualitatively proportional to the latter, and since production can be changed by assigning a worker at a city level, city population seems to be a better choice. This kind of reasoning is a novel use of qualitative representations that straddles the boundary of QR and language understanding. Given this conclusion, the qualitative statement that would be extracted is

```
(qprop (Production Boston)
       (Population Boston))
```

Suppose instead the previous sentence were something more abstract, e.g. "Consider the production in a city." Then

```
Q3:
  quantityType: Production
  quantityEntity: City
```

and the coreference merge would combine Q3 and Q1. Assuming the same kind of reasoning occurred, the qualitative statement that would be extracted is

```
(qprop+TypeType Production Population
                City City same)
```

That is, a city's production depends in part on its population. Thus the same frame interpretation should be able to generate both type level and instance-level representations, depending on context.

The second example is about heat flow, from a book on solar energy [2], i.e. "The heat flows from the brick to the ground, because the brick has a higher temperature than the ground." This compact sentence specifies an instance of heat flow. The verb "flows" denotes a process, in this case HeatFlowProcess, which is already in the knowledge base, as is ThermalEnergy for "heat". The prepositions "from" and "to" indicate direct influences [19]. So the model fragment frame would be

```
Process1
     participantOf: {Brick1, Ground1}
     conditionOf: Ord1
     consequenceOf:{Dinfluence1,
                    Dinfluence2}
     statusOf: Active ;; Because the flow is occurring
```

```
Dinfluence1
 constrained: ((QPQuantityFn ThermalEnergy)
               Brick1)
 constrainer: (RateFn Process1)
 influenceType: Direct
 influenceSign: -
 sourceOf: Process1

Dinfluence2
 constrained: ((QPQuantityFn ThermalEnergy)
                Ground1)
 constrainer: (RateFn Process1)
 influenceType: Direct
 influenceSign: +
sourceOf: Process1
```

We have written the constrainer and constrained arguments to the influences in assertional form rather than frame form to save space.

The causal connection will be drawn between an ordinal and the active status of the process, i.e.

```
(causes-FrameProp Ordinal1
       (statusOf Process1 Active))
```

where

```
Ordinal1
 quantity1: ((QPQuantityFn Temperature)
             Brick1)
 quantity2: ((QPQuantityFn Temperature)
             Ground1)
 ordinalReln: greaterThan
```

Again, we are using assertional forms for the quantities to save space. The ontological status of entities in examples (e.g. Brick1 and Ground1 here) in an interesting question. An instance-level model would be the mostly likely default for a reader, given that the original text includes specific diagrams to go with the text, and uses analogies to tie these ideas to more general phenomena. Generalizing to the specific categories, i.e. all bricks and the ground, would also be reasonable, in which case either a logically quantified description or a type-level description would be generated.

5 Progress on Implementation

While we are far from a full implementation at this writing (8/20/20), we have made some progress. Detection criterion for a broader variety of quantity frames has been implemented, exploiting the richness of the OpenCyc ontology. For example,

"rough carpet"

yields a frame whose quantityType is SurfaceSmoothness and whose quantityValue is the symbol Rough, which is part of a scale of values linked via ordinal relations. Similarly,

"John is heavy"

yields a frame whose quantityType is Mass and whose quantityValue is (HighToVeryHighAmountFn Mass).

In some cases, multiple interpretations are produced, e.g.

"John is fast"

yields one interpretation with the `quantityType` being `Speed` and `quantityValue` being `(HighAmountFn Speed)`, while another has the `quantityType` being `Time-Quantity` and `quantityValue` being `ShortTime`, another OpenCyc symbolic value. Ambiguities are handled either by abductive preferences or by semantic filtering, maintained so that higher levels of interpretation have material to work with.

Ordinal frames exploit the semantics of comparatives, which explicitly identify the type of quantity as well as the ordinal relation, e.g. "The elephant is heavier than the fly" introduces quantity frames for `Mass` for both the elephant and the fly, which are used as slot-fillers for the ordinal frame generated.

The example sentence "Production depends on population." yields an influence frame which is `Qprop` and `+`, with the `quantityTypes` for the two quantity frames being `Production-Freeciv` and `(MeasurableQuantityFn cityPopulation)` respectively. The recognition of influences is simplified by arranging the detection rules in a priority ordering, so that the quantities are recognized first, and hence causal connections involving discourse variables representing the quantities can be used to infer the qualitative proportionality. Another use of priority ordering is within quantity detection itself – rules which detect quantity frames with only quantity types, e.g. used in this example, are only run after those which detect quantity frames with additional information, such as entities and values. A non-monotonic test is used in the higher-level detection rules to skip adding an abstract frame when a more concrete frame has already been constructed to explain the same piece of text.

6 Discussion

The design discussed here draws on the best properties of prior systems for QP frames, potentially providing a unified system that can handle the automatic construction of both instance-level and type-level qualitative representations as appropriate, based on context and world knowledge. Our next step is to finish implementing narrative functions for automatically introducing the full set of unified QP frames as part of the Companion NLU system and algorithms for automatically formulating both instance-level and type-level QP models as needed to support reasoning. We plan to experiment with this new system over both our existing corpus of texts and on the new challenge problems being formulated by AI2 and others. Our goal is to demonstrate that high-precision understanding that construct explicit qualitative representations can produce more explainable results, with higher data-efficiencies [4] than alternate approaches.

ACKNOWLEDGEMENTS

This research was sponsored by the Air Force Office of Scientific Research and the Office of Naval Research.

REFERENCES

[1] Allen, J. (1995). *Natural language understanding, 2nd Edition*. Redwood City, CA: Benjamin/Cummings.

[2] Buckley, S., (1979). *From Sun Up to Sun Down*. New York: McGraw-Hill.

[3] B. Bredeweg, A. Bouwer, J. Jellema, D. Bertels, F. Linnebank, and J. Liem. 2007. Garp3 - A new workbench for qualitative reasoning and modelling. *4th International Conference on Knowledge Capture (K-CAP 2007)*, D. Sleeman and K. Barker (eds), pages 183-184, Whistler, BC, Canada, October 2007.

[4] Crouse, M., MFate, C.J., and Forbus, K.D. (2018). Learning from Unannotated QA Pairs to Analogically Disambiguate and Answer Questions. *Proceedings of AAAI 2018*.

[5] Fillmore, C. J., Wooters, C., & Baker, C. F. (2001). Building a Large Lexical Databank Which Provides Deep Semantics. *Proceedings of PACLIC-2001*.

[6] Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85-168.

[7] Forbus, K. and Gentner, D. (1997). Qualitative mental models: Simulations or memories? *Proceedings of the Eleventh International Workshop on Qualitative Reasoning*, Cortona, Italy, June 3-6, pp. 97-104.

[8] Forbus, K. & Hinrichs, T. (2017) Analogy and Qualitative Representations in the Companion Cognitive Architecture. *AI Magazine*, 38:34-42.

[9] Grice, Paul (1975). "Logic and conversation". In Cole, P.; Morgan, J. (eds.). *Syntax and semantics. 3: Speech acts*. New York: Academic Press. pp. 41–58.

[10] Hayes, P. (1981). The Logic of Frames. In *Readings in Artificial Intelligence*, Morgan-Kaufmann.

[11] Hinrichs, T. and Forbus, K. (2012). Toward Higher-Order Qualitative Representations. *Proceedings of the Twenty-sixth International Workshop on Qualitative Reasoning*. Los Angeles CA, July.

[12] Jia, Robin, and Percy Liang. "Adversarial Examples for Evaluating Reading Comprehension Systems." Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing. 2017.

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

[14] Klenk, M., Forbus, K., Tomai, E., Kim,H., and Kyckelhahn, B. (2005). Solving Everyday Physical Reasoning Problems by Analogy using Sketches. *Proceedings of 20th National Conference on Artificial Intelligence (AAAI-05),* Pittsburgh, PA.

[15] Kuehne, S. and Forbus, K. (2004). Capturing QP-relevant information from natural language text. *Proceedings of the 18th International Qualitative Reasoning Workshop*, Evanston, Illinois, August

[16] Marcus, Gary. "Deep learning: A critical appraisal." arXiv preprint arXiv:1801.00631 (2018).

[17] C. Matuszek, J. Cabral, M. Witbrock, and J. DeOliveira. 2006. An Introduction to the Syntax and Content of Cyc. In 2006 *AAAI Spring Symposium on Formalizing and Compiling Background Knowledge and Its Applications to Knowledge Representation and Question Answering*, (Palo Alto, CA).

[18] McFate, C.J., Forbus, K. and Hinrichs, T. (2014). Using Narrative Function to Extract Qualitative Information from Natural Language Texts. *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, Québec City, Québec, Canada.

[19] McFate, C.J. and Forbus, K. (2015). Frame Semantics of Continuous Processes. *Proceedings of the 28th International Workshop on Qualitative Reasoning (QR2015).* Minneapolis, MN.

[20] McFate, C., and Forbus, K. (2016). Scaling up Linguistic Processing of Qualitative Process Interpretation. *Proceedings of the Fourth Annual Conference on Advances in Cognitive Systems.* Evanston, IL.

[21] Minsky, M. (1975). A Framework for Representing Knowledge. In P. Winston (Ed.), *The Psychology of Computer Vision*. McGraw-Hill.

[22] Tafjord, Oyvind, Peter Clark, Matt Gardner, Wen-tau Yih, and Ashish Sabharwal. "Quarel: A dataset and models for answering questions about qualitative relationships." In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 7063-7071. 2019.

[23] 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.

[24] Weld, D. (1990). *Theories of Comparative Analysis*. MIT Press.