# Qualitative Reasoning for Learning by Reading: A Theoretical Analysis

**Kenneth D. Forbus, David Barbella, Clifton McFate**

Qualitative Reasoning Group, Northwestern University

2133 Sheridan Road, Evanston, IL, 60208, USA

forbus@northwestern.edu

### Abstract

One of the original motivations for qualitative reasoning was to capture the informal, intuitive notions about the continuous world that we all share, learned via a combination of experience and culture. For example, prior research suggests that qualitative dynamics can play an important role in natural language semantics. However, the constraints of everyday qualitative reasoning are different from more technical professional reasoning contexts, such as engineering. This paper examines qualitative reasoning in the context of learning qualitative dynamics for domains via reading texts. Based on experience with generating qualitative models from texts, we argue that qualitative reasoning for such everyday models raises new problems for qualitative reasoning, which opens up new research frontiers.

## Introduction

Qualitative reasoning was intended to capture both the intuitive, everyday models of the person on the street and the more rigorous models that underlie scientific and engineering knowledge about the continuous world. As a field, we have had great success with modeling professional qualitative reasoning (e.g. de Kleer, 1984; Bredeweg et al. 2009), but much less energy has been put into investigating everyday qualitative reasoning. There are good reasons to believe that the requirements of everyday QR differ from the requirements of capturing expert scientific and engineering qualitative reasoning. The traditional model of doing qualitative reasoning involves domain theories that are complete relative to the phenomena to be modeled, often incorporating multiple levels of granularity and multiple perspectives. Reasoning is performed via automatic model formulation, using first-principles knowledge combined with modeling assumptions based on domain principles and experience. Qualitative simulation over complete qualitative states provides a mechanism for ensuring that all important categories of possible behaviors are generated, constructing possibilities which can then be explored via more detailed knowledge as necessary (de Kleer & Brown, 1984; Kuipers 1994, Bredeweg et al. 2009). People just learning a domain are different. Their knowledge is partial, both in terms of what processes and phenomena are relevant in a domain, but also what they know about the phenomena that they have learned about. Despite this, they are still able to reason through quite complex situations, often extending their knowledge as they do so.

This paper is a theoretical investigation on what is required to extract useful qualitative representations by reading text. It brings together two lines of QR research. The first is examining the role that QR plays in natural language semantics (Kuehne 2004; McFate et al. 2014). As outlined below, this has led to the ability to extract most constructs of QP theory from simplified English text. What it does not do is provide an account of how such fragmentary, partial knowledge can be used subsequently to do qualitative reasoning, including looking for misconceptions of the type that are inevitable given the nature of natural language as a communication channel. The second line of QR research has examined how analogy and similarity provide a family of alternate methods for performing qualitative reasoning (Forbus & Gentner, 1997; Yan & Forbus, 2004; Friedman, 2012). As outlined below, these techniques rely on having a library of previous experiences, which can be used by analogy to explain new situations, but also to produce more rule-like knowledge as experience accumulates. These lines of investigation provide some, but not all, of what is needed to learn useful qualitative knowledge via reading. This paper walks through the entire process of learning by reading and using the learned knowledge, to identify the gaps and open problems that need to be addressed.

We begin by providing some brief relevant background about prior work, to set the stage. Then we walk through five tasks involved in learning and using qualitative knowledge from reading, to see where existing ideas probably suffice and where new ideas will be needed. We close with some conclusions and a discussion of future work.

## Background

We start by summarizing the models of analogical processing we build upon, then discuss Friedman's model of conceptual change, which provides an excellent starting point for modeling everyday qualitative reasoning.

We build on Gentner's (1983) structure-mapping theory of analogy, which describes analogy and similarity in terms of computing comparisons between structured, relational representations. Our work on learning by reading uses models of three analogical processes. SME (Falkenhainer et al. 1989) models analogical matching. In learning by reading, for example, SME is used to process instructional analogies and to construct suggestions for word sense disambiguation by analogy with previous choices. MAC/FAC (Forbus et al. 1995) models similarity-based retrieval. In learning by reading, it is used to retrieve potentially relevant cases for generating questions about new material and to retrieve cases for word sense disambiguation. SAGE (McLure et al. 2010) models analogical generalization. In learning by reading, SAGE is used to construct more portable knowledge about disambiguation choices and to build up models about concepts described in texts.

Friedman's (2012) *assembled coherence theory*, and the TIMBER implementation of it, uses these three analogical processes and QP theory to model human conceptual change. Friedman proposes that people store local explanations of phenomena, either acquired experientially or culturally, where the reasons for behavior are ultimately grounded out in first-principles model fragments. Explaining a new behavior involves using MAC/FAC to retrieve a prior explanation, and then applying the model fragments from that explanation to the new situation, using an abductive process to make missing assumptions as needed. Preference criteria take the cost of an explanation into account, i.e. assuming an unknown condition is far less costly than living with a contradiction. Finding lower-cost models drives the process of conceptual change. Implicit in this model is the idea that a process like this forms the basis for everyday qualitative reasoning, and we agree that this is very plausible, especially with two extensions proposed below.

## Five Tasks in Learning by Reading

We include using learned knowledge as part of the process of learning by reading, since a reader's understanding is never completely accurate, and hence trying out what was gleaned is an essential part of refining the knowledge into something useful. We will use examples drawn from two chapters from a book on solar energy, *Sun Up to Sun Down* (Buckley 1979), intended for non-specialists. Chapter 2 concerns the difference between heat and temperature and the basics of heat flow, which are explained using an analogy between water and heat. Chapter 16 concerns the operation of a solar hot water heating system, part of an extended multi-chapter example that tracks the operation of the system through a typical day. The five tasks we use in this analysis are (1) *achieving initial understanding*, (2) *knowledge integration*, (3) *question-answering from partial models,* (4) *similarity-based qualitative simulation*, and (5) *detecting, diagnosing, and repairing misconceptions*. We discuss each in turn.

```
(isa FluidFlow-Translation15256 QPProcessType)
(mfTypeParticipant FluidFlow-Translation15256 ?pan
     CookingVessel to-UnderspecifiedLocation)
(mfTypeParticipant FluidFlow-Translation15256 ?stove
     CookingRange from-UnderspecifiedLocation)
(mfTypeConsequence FluidFlow-Translation15256
     (i+ ((QPQuantityFn ThermalEnergy) ?pan)
       (RateFn ?self)))
(mfTypeConsequence FluidFlow-Translation15256
   (i- ((QPQuantityFn ThermalEnergy) ?stove)
     (RateFn ?self)))
```

Figure 1: Semantic interpretation for the sentence "Heat flows from the hot stove to the cool pan"

**Achieving Initial Understanding** This involves building up a semantic interpretation of the text. The EA NLU system (Tomai & Forbus, 2009) uses a chart parser and custom grammar, combined with ResearchCyc knowledge base contents. The semantic interpretation system is organized around Discourse Representation Theory (Kamp & Reyle, 1993), which provides methods for handling logical and numerical quantification, counterfactuals, and other useful semantic distinctions. The recognition of QP theory constructs is performed via *narrative functions* (McFate at el 2014), i.e. extracting the functional role of each sentence in the ongoing discourse. The semantic interpretation process is abductive, with the system preferring explanations of the text that justify consistent higher-order narrative functions.

Consider for example the sentence whose (partial) semantic interpretation is shown in Figure 1. The prepositional phrases (e.g. "to" and "from") in the context of a motion verb like *flow* provides the participants needed to support a direct influence reading. Getting to this interpretation is complex, because natural language is inherently ambiguous: It relies on context and on the prior knowledge of smart beings to interpret what is being said. Ambiguities arise in several ways. First, there can be multiple parses, involving different types of attachment. For example, in "The man saw a dog with binoculars," it is ambiguous whether the man or the dog has the binoculars. Second, there are typically multiple word senses for each word, e.g. "hot" can mean high in temperature, sexy, or spicy. The analysis constructed by the language system contains all of these possibilities and axioms linking their entailments together, so that the abductive process can make choices that lead to producing QP knowledge. However, there can be multiple sets of choices that produce QP knowledge, and not all choice sets are constrained by abduction. Other methods the system uses for resolving ambiguities include experience (i.e. analogical word sense disambiguation (Barbella & Forbus, 2013)) and domain-independent heuristics (e.g. prefer the interpretation that provides the most information).

One subtlety in this interpretation process is that some language statements are about things in general (i.e. generics) versus specific situations. Recognizing generics correctly is itself a difficult problem. Here the sentence is interpreted as a generic, because the system introduced logical variables for the participant roles.

People are somewhat capable of keeping track of the sources of their knowledge, more so if it is something that was recently learned. In Companions (Forbus et al. 2009), the knowledge gleaned from a particular reading of a text is stored in a set of linked cases, i.e. Cyc-style microtheories. This enables the knowledge to be easily used in subsequent reasoning, by including those microtheories in the logical environment of the reasoning, and to compare/contrast what is obtained from different sources, or the same source read at different times with different states of prior knowledge.

Since semantic interpretation is not deductively valid, mistakes can be made. This makes the next task tricky.

### Integrating Newly Learned Knowledge

Few of us start from nothing when reading, which means that what one reads must be integrated with what one already knows. In people, there are strong individual differences in how much this happens. Some appear to passively store the information from a text in a manner where it gets regurgitated on tests, but has minimal interaction with anything else that they know. Others appear to aggressively look for incompatibilities between the new and the old, and ask themselves how they can use the new knowledge. The process of *rumination* in Learning Reader (Forbus et al. 2007) provides one model of more aggressive processing. In rumination, the system asks itself questions. Learning Reader was focused on learning about world history, so it asked itself two kinds of questions. The first were basically forms of the Journalist's Questions about events (i.e. who, what, when, where, why, how), fleshing out what one typically knows about events. The second were generated via analogy with similar types of entities, e.g. if what was read about was a military operation, it used MAC/FAC to retrieve the most similar prior operations it already knew about, and used the candidate inferences generated from the mapping(s) as queries about the new information.

Different question generation strategies are needed for QR knowledge. For elaboration of knowledge, questions aimed at filling out what is typically known about model fragments seem important. Processes can be introduced in text without describing their conditions or consequences, for example. Similarly, if a process has been introduced, but no constraints have been put on its rate, asking what it depends upon is a useful question. Moreover, processes are often introduced via concrete examples, e.g. "Heat flows from the Sun to the Earth."

This leads to an interesting question: What are the participants for this process, in general? Aside from idealizations explicitly introduced to help guide the application of ideas (e.g. point mass, infinite sink), such information is often left implicit in texts. This is not unreasonable, since it is far from clear that the upper levels of human ontologies are uniform across people. For example, Chapter 2 also communicates the relative nature of heat flow by using the example of a frozen bird in an oven. Here is the simplified English form that our system can process:

> Consider a situation. A frozen chicken leaves a freezer. The frozen chicken is placed into a refrigerator. The chicken warms. The chicken is in the cold refrigerator, but the heat flows into the chicken. Because the chicken is colder than the refrigerator, the heat flows into the chicken. This causes the chicken to warm.

One possible way to proceed is to look at the superordinate concepts of all of the concrete concepts mentioned, to see if there are pre-existing concepts that capture what is relevant about participating in heat flows. This does not always work. Consider for example the various concepts that are used in Chapter 2 as sources and destinations for heat flow. Using the ResearchCyc knowledge base and starting from `Oven`, `Brick`, `Chicken`, and `Ground`, there are twenty-three superordinate concepts shared between them. Unfortunately, they are all very generic, e.g. `Thing`, `PartiallyTangible`, `SpatialThing-Localized`. None of them are plausible candidates, since they don't capture what is key about the participants: That they are the kind of thing that can be modeled as having heat and temperature.

There are at least two strategies for handling participant constraints more generally, and especially in the case where there isn't a natural superordinate. The first is to use analogy in modeling, checking to see if a potential participant is sufficiently similar to one of the prior known occurrences of that process, especially by using analogical generalization (Klenk et al. 2008). The second is to introduce a specific concept to be the constraint used on each type of participant, and install inheritance links from concepts observed to be participants to that superordinate. Thus over time, the concept will become elaborated as the set of inheritance links grows.

Consistency checking is more open-ended. Additional knowledge about a process or a phenomena captured by a model fragment needs to be checked to see if they can be combined together. (Unlike the usual human-readable syntax for QP models, model fragments are represented by collections of individual assertions, making their dynamic combination feasible.) Non-local constraints (e.g. that a quantity can never be both directly and indirectly influenced, for example) could either be tested via a static analysis or by looking for such problems when doing future modeling tasks.

**Question answering by reasoning over partial models**

In our experience, the vast majority of questions that arise in popular science books concern within-state qualitative reasoning, so we begin with that, and treat prediction and *postdiction* (i.e. explaining how a state might have come about) separately below. Suppose we have a problem from a science test, such as

> When a person's sweat evaporates, the person feels cooler. Which of the following statements best describes why sweating helps the person feel cool?
>
> A. Heat is absorbed by sweat when it evaporates.
>
> B. Heat is absorbed by the body when sweat evaporates.
>
> C. The temperature of the water in sweat goes down when it evaporates.
>
> D. The temperature of the water in the body goes up when sweat evaporates.

Such problems typically include a scenario, introducing entities and relationships among them to be reasoned about using a combination of text and diagrams (here, a person sweating), and one or more questions about them. Extracting the meaning of the scenario involves the same reading process as reading the main text, except that the question about whether something is a generic versus specific statement is much more likely to be resolved in favor of interpretations involving specific entities. Questions must be translated into queries that can be operationalized in terms of qualitative reasoning operations (or other operations – most books involve multiple types of knowledge and intermingle them as needed). We view the division of labor involved in question understanding to be more a matter of reasoning than of language understanding. That is, the language system produces expressions involving verbs and other reasonably abstract relationships, which are then decoded into a sequence of QR operations by problem-solving methods operating over the output of the language analysis.

Unfortunately, the partial nature of the models constructed by language means that there often isn't enough information to conduct first-principles reasoning. Take the way heat flow is described in Chapter 2. It is described entirely of specific, concrete situations. The text provides the specific types of participants for particular instances of this process, but it does not provide information about process types directly. This is where similarity-based qualitative reasoning becomes crucial. Explanations of physical phenomena read in the text provide analogs that can be used to construct qualitative models for new situations, by mapping qualitative representations onto the new situation. Often not all of this knowledge can be mapped to the new problem: If the analog involves a situation where an object is getting colder, but there are other processes involved in the new situation (e.g. one or more inflows as well as outflow), the causal laws should transfer intact, but their conclusions, which are based on closed world assumptions that do not hold in the target situation, should not be.

How should relevant analogs be found? TIMBER uses a case library of prior explanations with MAC/FAC to retrieve explanations, but it also includes pointers between cases to indicate when an explanation has been superseded by a better one. This seems like a psychologically plausible approach. It captures the fragmentary nature of most human mental models (e.g. Collins & Gentner, 1987), since different explanations for the same phenomena can be retrieved if the situations are dissimilar on the surface. Whether this will scale computationally is an open question at this time, given the difficulty of accumulating large-scale libraries of formally represented explanations. (This is one reason why we are researching learning by reading, because potentially a system can build up large bodies of knowledge by reading books and web pages.) But there is one problem with TIMBER for this task: It assumes that model fragments are fully specified in logically quantified terms. Model formulation, in TIMBER, is accomplished by retrieving a prior explanation and using the set of model fragments found in it as the domain theory for model formulation in the new situation. While the model formulation process was abductive, in that it would conjecture missing participants and preconditions (albeit at increasing cost to the explanation it was constructing), it still required first-principles model fragments. A different approach would be mapping model fragment instances via analogy, i.e. they would be candidate inferences from the retrieved situation. This has the advantage of not requiring the early introduction of logical variables, but the cost of having to do multiple mappings and/or rerepresentation if the number of entities is different. To see this, consider trying to explain a three container situation in terms of prior experience with a two-container situation (Figure 2). The prior experience (base) must be mapped twice, giving rise to conflicting predictions about the direction of change in the level of G. Re-running influence resolution on just the subset of the situation where the analogies cannot provide predictions seems like the most sensible approach, but this requires refactoring limit analysis to work in a more focused, local manner.

**Prediction and postdiction via analogy**

Generating multistate descriptions of behavior via analogy has been explored previously (Yan & Forbus, 2004). That model mostly used envisionments generated by a first-principles qualitative simulator (Gizmo), although one example was shown where a partial, hand-generated explanation was mapped to provide predictions about a new situation via analogy, which is encouraging. But is this the best that can be done? People are capable of estimating that some state transitions are more likely than others, for example, and that cannot be extracted from a purely first principles qualitative model.

Here is a proposal for a method of similarity-based qualitative simulation that provides estimates of likelihoods for transitions. Suppose incoming experience – and we include situations that the system reads about, as well as any physical experiences that it might have – is carved up into triples of the form

*<Before State, Transition, After State>*

That is, the propositional content of both the before and after state are part of the same case, and statements linking the particular cause of the transition (either a limit hypothesis occurring or a discrete action taken that changes a precondition) are also part of the same case. Suppose further that these triples are assimilated by SAGE into a generalization context. Then predictions for what happens next in a new situation can be made by (a) retrieving triples from SAGE, restricting the before state to match the before state in anything retrieved[1], and (b) using the frequency information associated with each transition (i.e., the frequency information that SAGE tracks for every statement in a case) to compute an estimated probability for each transition. There are a number of subtleties to be worked out, e.g. multiple generalizations could be formed that have analogous transitions, and these would have to be merged to maximize the accuracy of the computed probabilities. But this technique does have the potential advantage of maximizing the reuse of each experienced transitions.

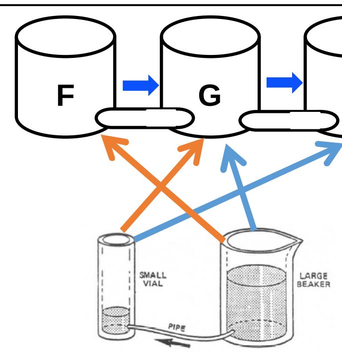


Figure 2: Explaining liquid flow in three containers in terms of prior experience with two containers requires multiple mappings (shown in red versus blue) and recomputing influence resolution for the properties of the liquid in G

[1] This can be done via *partition constraints* provided as part of the process. These essentially forbid matches between objects of different types.

**Detecting, diagnosing, and repairing misconceptions**

Errors in learned knowledge are inevitable. Such mistakes can be detected in several ways: A person might provide feedback about particular conclusions and/or reasoning steps, and the system itself might re-try prior problems in light of new knowledge to see if it now computes different answers than what was previously viewed as correct answers. Ideally, learned knowledge is used soon after acquired, to simplify the diagnosis process. We believe the approach of de Koning et al. (2000) for diagnosing problems in student models could be adapted for self-modeling in Companions in order to debug learned knowledge. They reified the reasoning done by a qualitative reasoner as a "device", whose components were mental operations (e.g. combining effects, retrieving specific facts) with information dependence between computations being represented as "wires" connecting the components. Given a discrepancy between a students' prediction for a situation and the value computed by this "device", they used the GDE algorithm (de KIeer & Williams, 1987) to diagnose possible reasons for the failure. For self-modeling in Companions, we think the following adaptation will enable this method to be used as-is: Track the frequency with which a particular fact, and/or the microtheory containing that fact, contribute to correct versus incorrect reasoning. This provides the probability information needed to rank candidate hypotheses. (One heuristic might be to assume that knowledge that has been around longer is more likely to be correct, assuming testing of it is reasonable uniform.) Another opportunity to detect problematic knowledge is to compare cases of learned knowledge about the same processes, either from the same source or from multiple texts, to look for discrepancies and opportunities to merge partial models.

## Conclusions and Future Work

We have argued that qualitative representations learned by reading place different demands on qualitative reasoning. Learned models are more fragmentary and more likely to have errors than hand-generated models. In some cases generic model fragments can be directly extracted from text, but even then, the information about participants supplied via language is more concrete than it should be for such model fragments to be widely applied.

Nevertheless, we believe that an important direction for qualitative modeling is to learn to work with such representations, generated by text, by dialogue, by experimentation in the physical world, and other sources of experience as AI systems become more connected to the world via high-bandwidth sensors (e.g. Kinect 2). Such reasoning, we have argued, needs to rely more heavily on analogy and be more local than traditional QR, which requires complete qualitative states. The TIMBER model, expanded with analogical model formulation and triple-based transition state generalization contexts, seems like a very promising approach.

Our next goal is to build up our techniques to the point where a Companion can read all of *Sun up to sun down*, and answer questions that we would reasonably expect people to answer after they had worked through that book. This is a difficult goal, but as the analysis here indicates, a novel combination of prior research results with a few extensions may indeed be enough to achieve this.


## Acknowledgements

This research was supported by the Intelligent Systems Program of the Office of Naval Research.


## References


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.

Barbella, D. and Forbus, K. (2013). Analogical Word Sense Disambiguation. *Advances in Cognitive Systems*, 2:297-315.

Bredeweg, B., Linnebank, F., Bouwer, A. & Liem, J. (2009) Garp3: Workbench for qualitative modeling and simulation. *Ecological Informatics* 4(5): 263-281.

Buckley, S. 1979. *Sun Up to Sun Down*. New York, NY: McGraw-Hill.

Collins, A. & Gentner, D. (1987). How people construct mental models. In D. Holland & N. Quinn (Eds.), *Cultural Models in Language and Thought*. pp.243-265, England: Cambridge University Press.

de Kleer, J. 1984. How circuits work. *Artificial Intelligence*, 24:205–280.

de Kleer, J. and Brown, J. 1984. A qualitative physics based on confluences. *Artificial Intelligence,* 24:7–83.

de Kleer, J. and Williams, B.C. (1987) Diagnosing multiple faults, *Artificial Intelligence* 32 (1) pp. 97-130

de Koning, K. Bredeweg, B., Breuker, J. and Wielinga, B. 2000. Model-Based Reasoning about Learner Behaviour. *Artificial Intelligence* 117(2):173-229.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial intelligence*, 41(1), 1-63.

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

Forbus, K. and Gentner, D. (1997). Qualitative mental models: Simulations or memories? *Proceedings of the Eleventh International Workshop on Qualitative Reasoning*, Cortona, Italy.

Forbus, K., Gentner, D. and Law, K. (1995). MAC/FAC: A model of Similarity-based Retrieval. *Cognitive Science*, 19(2), April-June, pp 141-205.

Forbus, K., Klenk, M., and Hinrichs, T. (2009). Companion Cognitive Systems: Design Goals and Lessons Learned So Far. *IEEE Intelligent Systems*, vol. 24, no. 4, pp. 36-46, July/August.

Forbus, K., Riesbeck, C., Birnbaum, L., Livingston, K., Sharma, A., and Ureel, L. (2007). Integrating Natural Language, Knowledge Representation and Reasoning, and Analogical Processing to Learn by Reading. *Proceedings of AAAI-07*, Vancouver, BC.

Friedman, S. E. (2012). Computational Conceptual Change: An Explanation-Based Approach. Doctoral dissertation, Northwestern University, Department of Electrical Engineering and Computer Science, Evanston, Illinois.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.

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

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

Klenk, M., Friedman, S., and Forbus, K. (2008). Learning Modeling Abstractions via Generalization. *Proceedings of the 22nd International Workshop on Qualitative Reasoning.*

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

Kuipers, B. 1994. *Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge*. MIT Press, Cambridge, MA.

McFate, C.J. and Forbus, K. (2011). NULEX: an open-license broad coverage lexicon. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*: Short Papers, 2, 363-367.

McFate, C.J., Forbus, K. and Hinrichs, T. (2014). Using Narrative Function to Extract Qualitative Information from Natural Language Texts. *Proceedings of AAAI 2014*, Québec City, Québec, Canada.

McLure, M., Friedman, S. & Forbus, K. (2010). Learning concepts from sketches via analogical generalization and near-misses. *Proceedings of CogSci10*

McLure, M.D., Friedman S.E. and Forbus, K.D. (2015). Extending Analogical Generalization with Near-Misses. *Proceedings of the AAAI 2015*, Austin, Texas.