Similarity-based Qualitative Simulation

Jin Yan (j-yan@northwestern.edu)
Kenneth D. Forbus (forbus@northwestern.edu)
Qualitative Reasoning Group, Northwestern University, 1890 Maple Avenue
Evanston, IL, 60201, USA

Abstract

People are remarkably good at using their common sense to predict and explain behavior. Qualitative modeling has provided formalisms that seem to capture many important aspects of human mental models, but standard qualitative simulation algorithms have properties that make them implausible candidates for modeling the flexibility, robustness, and speed of human reasoning. This paper describes a different approach, *similarity-based qualitative simulation*, which uses standard QR representations but with analogical processing to predict and explain behaviors. We summarize the problems with existing QR processing models as psychological accounts and outline the theory of similarity-based qualitative simulation. We illustrate the utility of this approach by a number of examples, generated with a prototype implementation.

Introduction

People are capable of using common sense knowledge to explain and predict everyday physical phenomena, such as filling a cup of tea, boiling a pot of water, kicking a pebble, or throwing a bowling ball. The models people use in reasoning about the physical world are called *mental models* [Gentner & Stevens, 1983]. Understanding mental models is a central problem in cognitive science, because they are crucial to understanding how everyday, common sense reasoning works. Qualitative reasoning research was originally motivated in part by the goal of creating a computational account of mental models [de Kleer & Brown, 1984; Forbus, 1984; Bredeweg & Schut, 1991]. Qualitative models do indeed capture several key features of mental model reasoning. These include representing partial and inexact knowledge, reasoning with partial knowledge, and generating multiple predictions at an abstract, conceptual level of representation. We believe that the representations developed by the QR community provide valuable formalisms for expressing the contents of human mental models. Psychological evidence supporting this comes from protocol studies [Bredeweg & Shut 1991; Forbus & Gentner 1986; Kuipers & Kassirer 1984] and studies of event perception suggest that people seem to comprehend motion in terms of discrete pieces [Hegarty, 1992; Zacks *et al.,* 2001].

However, we also see significant problems with qualitative simulation, as it has been typically defined in the QR community, when viewed as an account of human mental model reasoning. [Forbus & Gentner, 1997] described three key problems:

1. *Excessive branching*. A huge number of possible behaviors can often be generated even for relatively simple situations, and the number of behaviors tends to grow exponentially with the size of the system simulated [Kuipers, 1994]. This makes standard qualitative simulation algorithms problematic as psychological models in two ways. First, they produce more possible outcomes than people do when making predictions with the same information. Second, human mental models reasoning is typically quite fast, and seems to scale better.
2. *Spurious behaviors*. Many spurious behaviors tend to be included in predictions of today's qualitative simulators [Kuipers, 1994]. Such behaviors logically follow from the low-resolution input qualitative descriptions but are not physically possible. We have never seen a protocol or study where subjects spontaneously mention, for instance, ordinal relationships between higher-order derivatives, even though such relationships must be considered for accurate qualitative reasoning from first principles.
3. *Exclusive reliance on generic domain theories*. Generic domain theories are attractive because they enable a broad range of possible systems to be modeled. However, people seem to understand and reason about the physical world by relying more on concrete, specific knowledge.

Forbus & Gentner [1997] proposed that *hybrid qualitative simulation*, combining similarity-based reasoning with first-principles reasoning, would provide a more plausible psychological account of human mental models reasoning than traditional purely first-principles qualitative simulation. The idea is that most predictions are carried out via analogical reasoning, based on experience with similar situations. With enough experience, and accelerated via the use of language, more abstract principles are slowly formed by a conservative generalization process (see Section 2). These principles are also available for something closer to first-principles reasoning in qualitative simulation.

This paper describes an important exploration of this model. Just as many early investigations into qualitative reasoning focused on purely qualitative reasoning, in order to better understand what it could contribute, here we focus on using analogy exclusively for qualitative simulation, what we call *similarity-based qualitative simulation*. Section 2 briefly reviews the analogical processing ideas we are building upon, and Section 3 describes the theory of similarity-based qualitative simulation that we have developed. Section 4 describes our prototype, and Section 5 illustrates its operation on several examples, analyzing the strengths and weaknesses apparent so far in our approach. Section 6 discussed related work, and Section 7 summarizes and discusses of future work.

2. Similarity-based reasoning

Human reasoning appears to rely heavily on analogy and similarity [Gentner & Markman, 1997]. In artificial intelligence, this observation has led to important work on *case-based reasoning* (CBR) systems, where reasoning is based on remembering [Leake, 1996]. CBR systems retrieve the most relevant cases from memory and adapt them to meet the new situations instead of using purely first-principles reasoning [Kolodner, 1993; Leake, 1996]. Although CBR systems originally aimed to provide human-like performance, most of today's CBR systems tend to rely on feature vectors. Unfortunately, there is ample psychological evidence that human cognition centrally involves similarity computations over structured representations [Gentner & Markman, 1993].

Our theoretical framework of similarity-based reasoning is based on Gentner's [1983] *structure-mapping theory*, and the computational model is based on the Structure-Mapping Engine (SME) for comparison tasks [Falkenhainer *et al.*, 1989; Forbus *et al.*, 1994] and MAC/FAC [Forbus *et al.*, 1995] for retrieval tasks. Given two descriptions, a *base* and a *target*, SME computes one or two *mappings* representing structural alignments between them. Each mapping contains a set of *correspondences* that align particular items in the base with items in the target, *candidate inferences*, which are statements about the base that are hypothesized to hold in the target by virtue of these correspondences, and a *structural evaluation score*, which provides an indication of the quality of the match, based on structural properties. Candidate inferences can contain *analogy skolems,* entities hypothesized in the target because of statements in the base. (A historical example of such an entity is caloric, a fluid postulated by virtue of an early analogy between heat flow and water flow.) SME has been used to simulate comparison processes and their roles in various cognitive processes. Here, we use SME to match previously stored behaviors to new situations, generating predictions by projecting the correspondences through state transitions predicted via candidate inferences.

MAC/FAC models similarity-based retrieval as a two-stage process. The first stage (MAC) uses a cheap, nonstructural matcher to quickly filter potentially relevant items from a pool of such items. These potential matches are then processed in the FAC stage by a more powerful structural matcher, namely SME. Thus the output of MAC/FAC is a set of mappings, where the targets are the probe and the bases are the retrieved memory items.

In addition to matching and retrieval, we believe that *generalization* over experiences has an important role to play in hybrid simulation. SEQL [Skorstad *et al.*, 1988; Kuehne *et al.*, 2000] models this generalization process through *progressive alignment*, using SME to compare examples incrementally and build up new generalizations by keeping the overlap when there are very close matches. However, at this stage of our investigation we have not incorporated SEQL into our system, so generalization will get little attention in this paper.

Another key process in analogy is *rerepresentation*, the process of changing the representations in sound ways to improve matching [Yan *et al.*, 2003]. There are three aspects to rerepresentation in our model: *detecting opportunities* for rerepresentation, *generating rerepresentation suggestions* based on libraries of general methods, and *strategies* for controlling the rerepresentation process. It works like this:

1. Opportunities for rerepresentation are detected using criteria based on the principles of structure-mapping theory (e.g., a "hole" in an argument, or many to one matches).
2. For each opportunity, *rerepresentation suggestions* that suggest ways to change the descriptions to improve the match are retrieved and tried.
3. One or more suggestions are adopted, causing changes in the base and/or target.
4. The match is re-performed with the updated base and target descriptions.
5. The process continues until the match is suitable, or it fails, as determined by the task's rerepresentation strategy.

Rerepresentation is important because it expands the space of situations for which each example behavior can be used, thereby improving the coverage they provide.

3. Similarity-based Qualitative Simulation

We propose *similarity-based qualitative simulation* as an alternative to the traditional purely first-principles approach typically used in QR. Similarity-based qualitative simulation relies on a library of remembered experiences and generalizations drawn from them and analogical processing to understand new situations. Specifically,

- *Prediction:* Given a new situation, similarity-based retrieval and analogical comparison is used to map a remembered physical behavior onto the situation. The predictions produced by these analogies, we conjecture, correspond to the content of mental simulations.
- *Abduction:* Given a behavior to be explained, an explanation is constructed by mapping explanations for remembered behaviors onto the new behavior.

In both cases, some first-principles reasoning may be used to help check analogical inferences, to filter aspects of the projected behaviors that do not make sense. But new behaviors are always generated via analogy, rather than via first-principles reasoning. This is how similarity-based qualitative simulation differs from the hybrid model proposed in [Forbus & Gentner, 1997].

The ability for the same analogical reasoning mechanisms to handle both within-domain and cross-domain analogies should provide a flexibility and smoothness to prediction and abduction that is more in accord with human behavior. Since multiple behaviors can be retrieved and applied, branching predictions are possible, just as they are with first-principles qualitative simulation.

SQS can exploit knowledge at different levels of abstraction. We assume that people start by accumulating specific memories, including many concrete details, such as visual appearances. Successive comparison of concrete situations leads to the formation of partial abstractions, which can be viewed as prototypical behaviors or *situated rules* [Forbus & Gentner, 1986]. In some circumstances, especially when aided by language and explicit instruction, situated rules become further generalized and augmented to form first-principles knowledge. Importantly, none of these types of knowledge replaces the other: We assume that all remain available, and should be usable in SQS.

Rerepresentation also plays an important role in SQS, since it extends the range of applicability of examples. Plus, the process of rerepresentation appears to change memory contents in ways that promote transfer [Gentner *et al.*, 2003].

4. A prototype SQS system

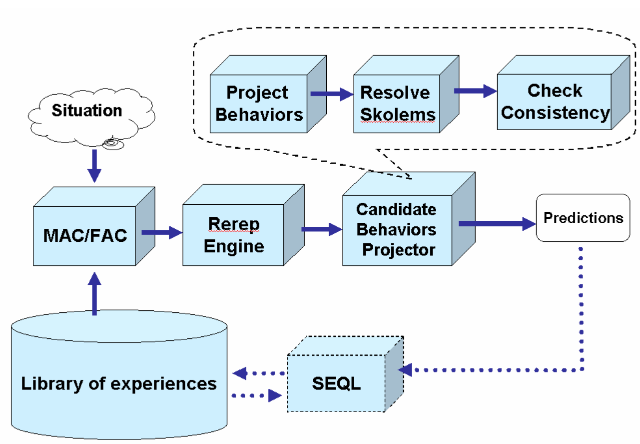

Figure 1: SQS system architecture

Figure 1 illustrates the architecture of our SQS system. The input is a situation, and the desired output is a prediction of the state (or states) that might happen next. Processing begins by using MAC/FAC on a library of experiences. MAC/FAC returns between zero and three remindings; if there is no reminding then no prediction is possible. If there are multiple remindings, the reminding with the highest structural evaluation score (i.e., the closest match to the situation) is selected for processing[1]. Second, the match between the retrieved situation and the current situation is scrutinized by the rerepresentation system, and tweaked if necessary to ensure that there are candidate inferences concerning state transitions, since these are what will provide predictions. Currently, all rerepresentation methods that might improve the match will be carried out, exhaustively. If rerepresentation fails, the system returns to the original match.

[1] This is a simplification; multiple retrievals are one way to generate multiple possible behaviors.

The third step is to use the correspondences and candidate inferences of the mapping to project possible next states. Let S be the initial situation, and Rs be the retrieved state mapped to it. Transitions from Rs to another state (say Rn) are part of the description of Rs, and since S has no transitions (by assumption), information about transitions will appear as candidate inferences. These inferences will contain an analogy skolem, representing "something like" Rn. SQS creates a new entity, say Sn, to represent the analog of Rn with respect to S, then retrieves facts from the experience library about Rn and projects them onto Sn by extending the mapping with a correspondence of Rn ↔ Sn.

The substitution process for generating new predictions is likely to lead to other analogy skolems (i.e. additional unknown objects conjectured for the target domain), which need to be resolved if possible. This means either identifying (or conjecturing) suitable entities in the target to be aligned with that base item. The conditions that the skolem must satisfy are extracted from the candidate inferences and solved for by a reasoning system. If no existing entity is found, then a new entity is created and the candidate inference constraints are applied to it.

Finally, each expression proposed about the target is checked for consistency and adapted if necessary. Two tests are used to determine consistency: (1) argument constraints associated each predicate are enforced and (2) each proposition should not be provably false [Falkenhainer 1988]. An alternate target correspondent will be sought when an inconsistency occurs. If the inconsistency cannot be resolved, the system returns to the next best reminding to restart the behavior projection process, until a consistent predicted behavior is formed for the current situation.

Our current prototype is still missing several important features. Currently, we carry out rerepresentation suggestions exhaustively; however, this process should be more selective and be controlled by task specific strategies. We neither store back into memory the results of rerepresentation, nor do we use SEQL to create generalizations on the fly. Even with these limitations, however, we think that the prototype shows some intriguing behaviors and possibilities. We next describe the prototype's operation on several examples, to illustrate its strengths and weaknesses.

5. Examples

To test the prototype, we generated a small library of experiences in two ways. First, we used Gizmo Mk2, a descendant of the original QP implementation, to generate envisionments for several classic QR examples (two containers, simple heat flow). We saved each state as a single case in MAC/FAC's case library, including its individuals, concrete details (e.g., individual appearance and/or surface properties), assumptions, ordinal relations involving both amounts and derivatives, model fragments, and transitions to possible next states, etc. Second, we generated by hand qualitative descriptions of behavior for a feedback system, to test the system's ability to work with behaviors involving incomplete state descriptions where no first-principles domain theory is available.

Example 1. Heat flow

Heat flow is a common phenomenon in physical systems. In order to provide behavioral predictions for the hot brick immersed in cold water situation (as shown in figure 2), the prototype begins by searching memory for analogous situations. Only one candidate analogue demonstrates strong similarities with the observed situation – heat flow from hot coffee to ice cube. Figure 2 demonstrates the behaviors of the hot coffee ice cube heat flow scenario, in which heat flows from the hot coffee to the ice cube, through a silver bar connecting them. Eventually a new state is reached in which the hot coffee and the ice cube have the same temperature, and the heat flow process has stopped. The prototype determines that the roles of the hot coffee, ice cube and bar in the heat flow description correspond to the roles of the brick, water and the surface contact between the brick and water in the target situation, respectively. It also generates candidate inferences that there should be a heat flow process active in the target input scenario, in which the temperature of the brick is dropping, while the opposite is true for the water. The projected new state for the hot brick cold water scenario is that the brick and the water reach the same temperature eventually, and heat flow process has stopped.

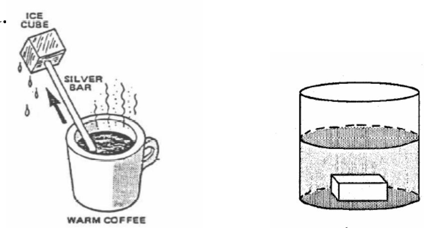

State0

```
↓ (Temperature Coffee)
↑ (Temperature IceCube)
(> (Temperature Coffee)
   (Tempearature IceCube))
(activeMF HeatFlow)
```

State1

```
→ (Temperature Coffee)
→ (Temperature IceCube)
(= (Temperature Coffee)
   (Tempearature IceCube))
(not (activeMF HeatFlow))
```

Input Scenario

```
↓ (Temperature Brick)
↑ (Temperature Water)
(> (Temperature Brick)
    (Tempearature Water))
(activeMF HeatFlow)
```

Behavior Prediction

```
→ (Temperature Brick)
→ (Temperature Water)
(= (Temperature Brick)
    (Tempearature Water))
(not (activeMF HeatFlow))
```

Figure 2. Similarity-based qualitative simulation for the hot brick immersed in cold water scenario

Example 2. Discrete on/off feedback control system

Consider a simple room heating system. Two major processes exist in this scenario. When the furnace is on, heat flows from the furnace to the room[2]. The room is also always losing heat to the outside. When the temperature in the room falls below the temperature set point on the dial, the furnace will turn on to heat the room. When the temperature in the house rises above the temperature set

Feedback Control System	**Water Level Regulation System**
Sensor	Floating ball
Comparator	Ball Stick
Controller	String + Pulleys
Actuator	Valve
Temperature set point	Proper water level
Room air	Tank water
Room	Water tank
Furnace	Water supply
Heat flow process	Liquid flow process
Furnace on process	Valve open process

Quantities	S1	S2	S3	S4	S5	S6
(Temperature Room) vs. SetPoint	<	=	>	>	=	<
(Ds (Temperature Room))		1			-1	
(activeMF FurnaceOn)		Yes			No	
(activeMF HeatFlow)		Yes			Yes	

Quantities	S1	S2	S3	S4	S5	S6
(Level TankWater) Vs. ProperWaterLevel	<	=	>	>	=	<
(Ds (Level TankWater))		1			-1	
(activeMF ValveOpen)		Yes			No	
(activeMF LiquidFlow)		Yes			Yes	

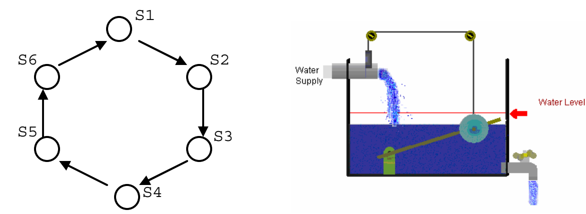

Figure 3. Similarity-based qualitative simulation for the water level regulation system

point, the furnace will turn off, and the continued heat loss from the room to the outside eventually causes the temperature of the room to drop below the set point. This behavior was encoded by hand, so that we did not have to make any particular commitments to the sources of lag or delay in the system, or any of a number of factors that would have been required in a correct first principles model of the situation. This enables us to see whether a partially understood situation might still be used for prediction. Figure 3 chart 2 summarizes the behavior for the heating system. Starting from any state, the behavior of the home heating system takes the path shown in the six states loop of Figure 3.

Now consider an analogous control system, also shown in Figure 3. The water tank system tries to keep the water level in the tank at a constant value when the faucet is turned on. There is a ball float connected to a stopper valve that moves up and down with the water level. When the ball float moves down, the stopper valve is open. When the ball float moves up, it lowers the stopper, closing the valve. The input scenario given to the prototype is a situation in which the ball float is lower than the proper water level, and the stopper valve is open and letting water into the tank from a larger water supply. The prototype maps distinct entities and processes between the water tank scenario and the home heating system respectively (as shown in the first chart of Figure 3). It retrieves state `S1` as the closest reminding for the input, and projects five subsequent future states, as shown the third chart of Figure 3.

Example 3. Proportional action control system

[2] We are ignoring how the furnace is kept hot, something every homeowner with a heating bill would like to do.

Proportional action control systems are different from discrete on/off feedback control systems (as modeled in Example 2) in that they provide correction proportional to the difference between the temperature and the set point. As the temperature approaches the set point, for instance, the output of the furnace is reduced. This reduces fluctuations in temperature.

When we give SQS a room heating scenario in which the temperature is lower than but approaching the set point and the controller is a proportional-action controller, it attempts to use the closest precedent in its memory. In fact, it will attempt to use the same memory it used to explain Example 2, i.e., the discrete-action controller behavior. The SQS prototype will thus predict that the furnace is fully on at its maximum rate, when actually its heat flow rate should be proportional to the difference between the temperature and the set point. This illustrates that similarity-based qualitative simulation can lead to spurious predictions, by applying inappropriate analogs. There are two points to make here. First, in a full hybrid simulator first-principles reasoning could be used to do additional testing of predictions, and could in some cases catch spurious behaviors if there is rich enough domain knowledge available. Second, similarity-based qualitative simulation can be improved by remembering the behaviors that it couldn't otherwise explain: A very simple, but we suspect very powerful, learning mechanism.

Example 4. Specific room heating system: a thermostat

This example illustrates how SQS applies knowledge from general schemas to make predictions for specific scenarios, and how the rerepresentation engine helps the prototype to achieve flexibility. As described in examples 2 and 3, the prototype's initial knowledge contains behaviors about the general schema for a room heating regulation system. Consider giving a specific room heating system, a thermostat, as the input scenario for the prototype, a small fragment of the representations involved in both situations is like:

B:
```
(senses SensorX (Temperature RoomAirX))
 (compares ComparatorX (Temperature RoomAirX)
           TemperatureSetpointX)
```

T:
```
(senses ThermostatY (Temperature RoomAirY))
(compares ThermostatY (Temperature RoomAirY)
           TemperatureSetpointY)
```

Our prototype determines that the thermostat plays the role of the sensor and comparator in the abstract schema. However, this match cannot be allowed, since it violates the 1:1 constraint of the structure mapping theory, which leads to only one pair of entity alignment (`SensorX` align to `ThermostatY`) being included in the legitimate match between the retrieved situation and the current scenario. [Yan *et al.*, 2003] calls such rerepresentation opportunities *rivals*, which are violations of the 1:1 constraint that lead to structural inconsistency of at least one match hypothesis. It is often caused by the same entity playing multiple roles in the same representation.

Next, our rerepresentation engine takes the reminding match that has the highest structural evaluation score, and refines the description of the thermostat, suggesting that it is the curvature of its bimetallic strip that measures the temperature, and the angular distance between the bimetallic strip and the dial's angle that provides the comparison. This is an example of an *entity splitting* strategy. In general, entity splitting strategies require identifying ways to divide up an entity into distinct parts or aspects, and rewrite its roles in the description to use one or the other of these parts or aspects. After rerepresentation, each of the aspects of the thermostat can match to distinct functional descriptions from the retrieved schema, leading to a much better match. E.g.,

T':
```
(senses (CurvatureFn BimetallicStrip)
        (Temperature RoomAirY))
(compares (AngleFn BimetallicStrip)
          (Temperature RoomAirY)
           TemperatureSetpointY)
```

Finally, the prototype uses the improved match to project new states for the rerepresented scenario.

6. Related Work

The observation that analogical matching could be used for deduction or abduction depending on the kind of knowledge involved was first made by Falkenhainer [1988]. His PHINEAS system used cross-domain analogies involving behaviors to first retrieve elements of a domain theory used to explain the original behavior, and then modify those domain theory elements to form a partial new domain theory that could explain the new behavior. This new domain theory was tested via first-principles qualitative simulation. While we build on ideas from PHINEAS in several ways, notably the use of the map/analyze cycle and using analogical reasoning for abduction, our MAC/FAC model for retrieval is more psychologically plausible than what was available then, and our focus is on using analogy directly for prediction and explanation, rather than constructing a new first-principles domain theory.

Connectionist simulations of analogical matching and retrieval, such as ACME [Holyoak & Thagard, 1989], LISA [Hummel & Holyoak, 1997] and CAB [Larkey & Love, 2003] are aimed at creating neurally plausible models of analogical processing. Unfortunately, such models so far cannot handle examples as complex as people can, including the examples described in this paper.

The CBR community has created many systems that generate predictions [cf. Kolodner, 1994]. However, CBR systems tend to use matching and retrieval systems that are optimized for each task and each domain. Our system uses psychological models of analogical matching and retrieval that have been used in a variety of domains and tasks [cf. Forbus *et al.*, 2002], making it more likely to scale to the breadth required for mental models reasoning.

7. Discussion & Future Work

We believe that similarity-based qualitative simulation is a plausible model of human reasoning. While our prototype is not yet complete, the examples suggest that the approach has promise. Most of the future work revolves around addressing its limitations:

- The library of experiences needs to be significantly expanded, to stress-test retrieval and rerepresentation.
- The rerepresentation process should be more selective and be controlled by task specific strategies.
- First-principles reasoning needs to be integrated with the existing SQS system to facilitate reasoning and filter possible candidate behaviors, turning SQS into a hybrid qualitative simulator.
- Learning strategies, in the form of storing back the results of rerepresentation and using SEQL to construct generalizations, also need to be explored.

Acknowledgements

This research is supported by the Office of Naval Research under the Cognitive Science and Artificial Intelligence Programs.

References

Bredeweg, B., & Schut, C. 1991. Cognitive plausibility of a conceptual framework for modeling problem solving expertise. *Proceedings of the 13th Annual Conference of the Cognitive Science Society,* 473-479. Hillsdale, New Jersey: Lawrence Erlbaum.

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

Falkenhainer, B., Forbus, K., Gentner, D. 1989. The Structure-Mapping Engine: Algorithm and examples. *Artificial Intelligence*, 41, pp 1-63.

Falkenhainer, B. 1988. Learning from physical analogies: a study in analogy and the explanation process. Report No. UIUCDCS-R-88-1479, Department of Computer Science, University of Illinois at Urbana-Champaign.

Forbus, K. 1984. Qualitative process theory. *Artificial Intelligence*, 24, 1984.

Forbus, K., Gentner, D. 1986. "Learning Physical Domains: Towards a Theoretical Framework" in Michalski, R., Carbonell, J. & Mitchell, T. *Machine Learning: An Artificial Intelligence Approach*, Volume 2, Tioga press.

Forbus, K., Ferguson, R. & Gentner, D. 1994. Incremental structure-mapping. *Proceedings of the Cognitive Science Society,* August.

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

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

Forbus, K., Mostek, T. & Ferguson, R. 2002. An analogy ontology for integrating analogical processing and first-principles reasoning. *Proceedings of IAAI-02*, July 2002.

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

Gentner, D. & Stevens, A. (Eds.) 1983. *Mental Models.* IEA Associates.

Gentner, D., Markman, A. B. 1993. Analogy - Watershed or Waterloo? Structural alignment and the development of connectionist models of cognition. In S. J. Hanson, J. D. Cowan, & C. L. Giles (Eds.), *Advances in neural information processing systems* (Vol. 5, pp. 855-862). San Mateo, CA: Morgan Kauffman.

Gentner, D., Markman, A. B. 1997. Structural alignment in analogy and similarity. *American Psychologist.*

Gentner, D., Loewenstein, J. & Thompson, L. 2003. Learning and Transfer: A general role for analogical encoding. *Journal of Educational Psychology*, 95(2), 393-408.

Hegarty, M. 1992. Mental animation: Inferring motion from static displays of mechanical systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 1084-1102.

Holyoak, K. J., Thagard, P. 1989. Analogical mapping by constraint satisfaction. *Cognitive Science*, 13(3), 295-355.

Hummel, J. E., Holyoak, K. J. 1997. LISA: A computational model of analogical inference and schema induction. *Psychological Review.*

Kolodner, J. L. 1994. *Case-based reasoning*. San Mateo, CA: Morgan Kaufmann Publishers.

Kuipers, B., Kassirer, J.P. 1984. Causal reasoning in medicine: analysis of a protocol. *Cognitive Science*, 8:363-385.

Kuipers, B. 1994. *Qualitative Reasoning: Modeling and simulation with incomplete knowledge.* Cambridge, Mass.: MIT Press.

Kuehne, S., Forbus, K., Gentner. D., & Quinn, B. 2000. SEQL: Category learning as progressive abstraction using structure mapping.

Larkey, L., Love, B. 2003. CAB: Connectionist Analogy Builder. *Cognitive Science*, 27, 781-794.

Leake, D. (Ed.) 1996. *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. MIT Press.

Ma, Joyce. 1999. A Case study of Student Reasoning about Feedback Control in a Computer-based Learning Environment. *29th ASEE/IEEE Frontiers in Education Conference*, San Juan, Puerto Rico

Skorstad, J., Gentner, D., & Medin, D. 1988. Abstraction processes during concept learning: A structural view. *In Proceedings of the Tenth Annual Conference of the Cognitive Science Society*. Montreal: Lawrence Erlbaum Associates.

Yan, J., Forbus, K., & Gentner, D. 2003. A theory of rerepresentation in analogical matching. In *The Proceedings of 25th Annual Conference of the Cognitive Science Society, August*

Zacks, J., Tversky, B. & Iyer, G. 2001. Perceiving, remembering, and communicating structure in events. *Journal of Experimental Psychology: General*, 130, 29-58