

Qualitative Mental Models: Simulations or Memories?

Kenneth D. Forbus

The Institute for the Learning Sciences

Dedre Gentner

Psychology Department

Northwestern University

1890 Maple Avenue, Evanston, IL, 60201, USA

Abstract

One of the original motivations for qualitative physics research was the creation of a computational account of mental models. For instance, a key intuition often associated with mental models is that they are *runnable*, i.e., there is a sense of deriving answers via mental simulation rather than logical reasoning. This paper examines three explanations for runnability, and argues that none of them is sufficient. Instead, a hybrid model combining aspects of all three is proposed, focusing on the integration of ideas from qualitative physics with ideas from analogical processing. Some psychological implications of this hybrid model are discussed.

Introduction

An active area of research in cognitive science is studying *mental models* (Gentner & Stevens 1983), the models people use in reasoning about the physical world¹. Understanding mental models is a central issue for cognitive science because they appear important in reasoning about complex physical systems, in making and articulating predictions about the world, and in discovering causal explanations for what happens around us. Mental models research also offers practical benefits. In an increasingly technological society, understanding the nature of mental models for complex physical systems and could help people learn better models which could reduce accidents and improve productivity (Norman, 1988).

One of the original motivations for qualitative physics research was to create a computational account of such mental models (de Kleer & Brown 1984; Forbus, 1984; Bredeweg & Schut 1991; White & Frederiksen 1990). A key intuition often associated with mental models is that they are *runnable*, i.e., there is a sense of deriving answers via mental simulation rather than logical reasoning. One explanation for runnability is that people are doing qualitative

simulation, based on general, first-principles knowledge of a physical domain. Another explanation is that we use some high-resolution mental simulation facility. Yet another explanation is that we are using memories of similar situations previously observed to construct a story about what will happen in the current situation.

In this paper we argue that none of these explanations by itself is adequate to account for human common sense reasoning. We believe a psychological account of qualitative reasoning will rely heavily on analogical reasoning in addition to reasoning from first principles. We propose a hybrid model, motivated by a combination of psychological findings and computational considerations. Aspects of this model have been tested by computer simulation. We believe that this hybrid account is more consistent with evidence about human learning than any of the pure models.

We begin by examining the three "pure" explanations for runnability. Then we outline our hybrid model and some of its psychological implications.

Three views of qualitative mental models

We focus on the use of mental models in common sense prediction tasks, like thinking about what might happen when filling a cup with coffee.

High-resolution mental simulations

An appealing intuition is that mental model reasoning is like watching a movie of a physical system with your mind's eye. This intuition has been the basis for proposals that link mental model reasoning with visual imagery (Funt, 1980; Kosslyn, 1980; Glasgow, 1992; Hegarty 1992; Gambardella, Gardin, & Meltzer, 1988). Evidence for our visual apparatus being involved in spatial reasoning includes psychological studies (Kosslyn, 1980), computational necessity arguments (Forbus, 1983; Forbus, Nielsen, & Faltings, 1991), and neuroscience experiments (Kosslyn, 1996). However, there is evidence that visual processing alone is insufficient. Hinton (1979) demonstrated distortions in mental imagery that do not fit within an array model of imagery, although they are consistent with mixed symbolic/metric representa-

¹ Our focus on physical domains and long-term knowledge structures distinguishes this sense of mental model from the other standard usage, e.g. Johnson-Laird (1983).

tions. Schwartz (1996) summarizes evidence suggesting that several kinds of knowledge are used in imagery, including physical, social, and haptic knowledge.

It seems likely that spatial mental models rely in part on visual computations. On the other hand, we know of no evidence suggesting that the data needed for quantitative simulation is available in common sense reasoning tasks, nor do we know of any evidence that people have a mental simulation facility capable of using such information. Consider predicting the pattern of liquid that will appear on a rug if a half-full cup of coffee is knocked off a table. Our visual apparatus is powerful enough to describe the shape of the patterns that result. However, we are not capable of predicting what specific shapes will result in advance. Solving this problem to a high degree of accuracy involves computational fluid dynamics; it seems quite unlikely that we are capable of performing such a prodigious feat mentally.

First-principles qualitative simulation

In some qualitative physics research, running a mental model is identified with carrying out a qualitative simulation of the system. Qualitative reasoning captures several important properties of mental model reasoning, namely

- *Handling incomplete and inexact data.* Qualitative information is easily extracted via perception, and such rough distinctions are more likely to be easily remembered than precise details.
- *Support for simple inferences.* Simple, everyday "obvious" inferences can be carried out easily. For instance, if nothing is happening, nothing is changing.
- *Representation of inexact knowledge.* Qualitative representations make causal knowledge explicit. They provide a vocabulary for expressing partial knowledge about causal theories and mathematical relationships, and methods to assemble this partial knowledge on demand for reasoning.
- *Representation of ambiguity.* In many everyday prediction tasks we can imagine several distinct outcomes. Qualitative simulations capture this ambiguity.

However, there are two problems with using current theories of qualitative simulation to account for mental model reasoning: excessive branching and exclusive reliance on generic models. We discuss each in turn.

Excessive branching Current qualitative simulators often produce a huge number of possible behaviors even for relatively simple situations (Kuipers, 1994). In some applications exploring every possible behavior, i.e., envisioning, is necessary (Shimomura, Tanigawa, Umeda, & Tomiyama, 1995; Price, Pugh, Wilson, & Snooke, 1995). But today's qualitative simulation algorithms tend to make many more distinctions

than necessary for most tasks, leading to unnecessary complexity in the behaviors they generate (deCoste, 1994). This makes them seem psychologically implausible, for two reasons². First, qualitative simulators often produce states containing many more distinctions than people report when considering the same scenario. Comparisons between derivatives of rates, for example, are needed for continuity calculations that rule out inappropriate state transitions. But we have never seen such comparisons mentioned in protocols. This does not by itself rule out their use internally. It could be that certain information is simply underreported in protocols. However, these discrepancies are grounds for asking whether such calculations are psychologically frequent. Second, the exponential nature of most qualitative simulation algorithms makes them implausible models for the rapidity of common sense reasoning.

Exclusive reliance on generic models The goal of most qualitative physics research is to build an idealized physical reasoner, a system that can reason with sophistication about the physical world in the way that the best human scientists and engineers do, without their frailties. This goal leads to a preference for systems that maximize generality and generativity. That is, the laws of qualitative physics are expressed in domain-independent terms, and knowledge of domains is expressed in situation-independent forms. It seems likely that people's mental models include laws and principles that are at least somewhat domain-independent, as well as domain knowledge that is situation-independent. But there is ample evidence suggesting that much of what people know about the physical world and how they reason about it is more concrete than that (Brown, Collins, & Duguid, 1989). The exclusive reliance of current qualitative simulation accounts on first-principles knowledge makes them implausible candidates for psychological models, except perhaps in very narrow ranges of high-expertise reasoning.

Memory-based reasoning

The third explanation is that running a mental model of a system corresponds to remembering how that system has behaved previously when in similar circumstances. The fact that people store and remember behaviors of physical systems is uncontroversial. How far memory-based explanations can go in explaining physical reasoning is still an open question. A major issue is generativity: How flexibly can past experiences be used to make new predictions, and especially predictions about novel systems and/or

² We do not know of experiments in the literature that address these questions: We are arguing based on our informal observations of people and qualitative simulators.

configurations? We believe there are three factors that make memory-based reasoning more generative than some might otherwise expect. First, qualitative representations reduce differences. Assuming people store and use qualitative representations of situations and behavior, then two situations that vary only in quantitative details will look identical with respect to the qualitative aspect of their behavior. Second, analogical reasoning can generate predictions for novel situations. For common sense reasoning, within-domain analogies (i.e., predicting what will happen when pouring coffee into a cup based on previous experiences pouring coffee into a different cup) should provide a reliable guide to action. Third, multiple analogies can be used to piece together models for complex systems (Spiro *et al* 1989).

There is psychological evidence that the same comparison processes used for cross-domain analogical thinking are also used for within-domain comparisons, in tasks ranging from visual perception to conceptual change (Gentner & Markman, 1997). It would be surprising if such processes were not used in common sense physical reasoning. However, memory-based reasoning alone is insufficient to explain our ability to use general-purpose, domain-independent physical knowledge – something that we undeniably do, even if there is disagreement over how much of it people do routinely and under what circumstances.

Similarity-based hybrid qualitative simulation

None of the pure models are sufficient to account for the runnability of mental models. We claim that a hybrid model is needed to explain the full range of human common sense reasoning. The model of reasoning and learning in physical domains we propose here differs from our previous model (Forbus & Gentner, 1987). We now suspect that the kinds of knowledge and processes that we previously divided into stages are actually more tightly interwoven. Specifically, we now believe that comparison processes play a central role throughout the span of expertise.

We begin by first examining what aspects we are adopting from each of three approaches discussed. Then we illustrate how predictions can be made in this hybrid model, providing a sense of “running” a mental model.

High-resolution mental simulations

We assume that some high-resolution representations are available for diagrammatic and spatial reasoning tasks, mainly through facilities shared with our visual systems. However, we assume that, spatial reasoning aside, there are no high-resolution mental simulations. Since assumptions about spatial reason-

ing are almost independent of the rest of the model, we ignore this issue in the rest of the paper.

What Qualitative Physics provides

We assume that people use many of the representational constructs of qualitative physics when reasoning about mental models. This includes

- Methods for representing partial information about numerical values, including signs (de Kleer & Brown, 1984), ordinals (Forbus, 1984), simple symbolic vocabularies (Guerrin, 1995), and order of magnitude relationships (Raiman, 1991; Mavrouniotis & Stephanopoulos, 1988).
- Causal and mathematical relationships capable of expressing partial knowledge (i.e., direct influences and qualitative proportionalities from QP theory and the extensions described in (Bobrow *et al* 1996)).
- Representations for modeling assumptions (Falkenhainer & Forbus, 1991; Rickel & Porter, 1994)
- Many of the ontologies that have been developed for specific domains, and multi-domain abstractions such as physical processes and devices.

We also assume that people encode varying amounts of detailed information about the values of continuous properties, in addition to qualitative properties.

We assume that people sometimes use domain-independent principles of qualitative reasoning and situation-independent general knowledge of particular domains. We also assume that much of people's physical knowledge is highly context-specific. That is, we assume that many principles of qualitative reasoning people use are domain-specific, and that much of their knowledge about a particular domain is also tied to situations or classes of situations within that domain. The difference between these may be seen by the following sequence of states of knowledge, each of which could be used for prediction, but takes quite different forms:

1. A remembered behavior concerning a specific cup at a specific time, e.g., more coffee pouring into your favorite cup leading to it flowing over the top and spilling on your desk. The behavior's description probably includes many concrete details, such as visual descriptions of the objects and their behaviors.
2. A remembered behavior concerning a specific cup at a specific time, including a causal attribution relating different factors or events, e.g., the overflow was caused by continuing to pour coffee once the cup was full. This attribution might come about by someone explaining the situation to you, or by analogy with an explanation given for another situation, or by the application of a more general abstraction. Additional qualitative relations might be included, such as blaming the overflow event on pouring a liquid, with the rate of overflow depending on the rate of pouring.

3. A generalization that coffee cups can overflow if you keep filling them up with liquid. This generalization might be formed by successive comparisons of very concrete situations, conservatively stripping away details that are not common across otherwise similar situations. Visual properties may be gone, but many aspects of the descriptions are still very concrete – coffee cups instead of containers, for instance, or even coffee instead of any liquid. More qualitative relationships may be included.

4. A generic domain theory of containers, liquids, and flow that supports limit analysis, e.g., the coffee cup is a container, the coffee in it is a contained liquid, therefore one limit point in the quantity space for the level of the contained liquid is the height of the cup's top, and that a qualitative transition in behavior will occur when the level (which is rising due to being the destination of a liquid flow, which is the only thing happening that is affecting the amount of coffee in the cup) reaches the height of the top of the cup.

The first state of knowledge represents pure memory. The last state of knowledge represents the sort of explanation that would be generated by first-principles qualitative simulators. They represent extremes on a continuum of knowledge about the physical world. The states in between represent what we suspect what might be very common in human mental models: intermediate levels of generalization and explanation, where partial explanations have been constructed in a conservative fashion (e.g., generalizing across liquids but still restricted to coffee cups). They are examples of what we could call *situated rules*, pieces of knowledge that are partially abstracted but still partially contextualized.

From an applications perspective, situated rules are the bane of good knowledge engineering practice. When engineering a domain theory, one strives for generality and broad coverage. In that context, the use of partially abstracted, partially contextualized knowledge represents a failure of analysis³. But the situations faced by knowledge engineers and by human learners are very different. Human learning is often very conservative, especially when someone knows little about a domain. Situated rules provide an intermediate form of knowledge between concrete/slightly schematized descriptions of behaviors and the mechanism-based ontologies of standard qualitative physics.

We conjecture that situated rules are used to express principles of qualitative physics as well as knowledge about particular domains. That is, it seems likely that there is a range of knowledge about physical reasoning, varying from concrete rules applicable to a small class of situations to the kinds of overarching, general

principles encoded in performance-oriented qualitative reasoning systems. English-speakers commonly use the phrase "what goes up must come down", and other language communities have similar expressions. How many of those speakers know that, assuming classical continuity, this statement implies the existence of an instant of time between going up and going down where the vertical velocity is zero? There is a large terrain between knowing nothing and having a broad-coverage general theory, and that terrain is not empty.

What analogical processing provides

Analogical processing provides several key capabilities:

Robust matching and inference. Structure-mapping theory (Gentner, 1983) provides an account of comparison processes and their roles in various cognitive processes that is consistent with a growing body of psychological evidence (Gentner & Markman, 1997). These computations have been simulated with SME (Falkenhainer *et al* 1989; Forbus *et al* 1994), which in turn has been used as a module in other simulations and in performance systems. Given two structured propositional representations as inputs, the base (about which more is presumably known) and the target, SME computes a *mapping* (or a handful of them). Each mapping contains a set of *correspondences* that align particular items in the base with items in the target, and *candidate inferences*, which are statements about the base that are hypothesized to hold in the target by virtue of these correspondences. SME can incrementally extend its mappings as more information is added to the base and/or target.

Integration of multiple types of knowledge. The same analogical processes can be used to operate on rules, concrete descriptions, and abstractions such as equations and plans (Forbus *et al* 1994). In the case of rules, the base is the rule and the target is the situation to which the rule is to be applied. Each mapping corresponds to an instantiation of the rule, with the candidate inference providing the new information⁴.

Incremental abstraction and rule generation. SEQL (Skorstad, Gentner, & Medin, 1988) uses SME in successive comparisons of examples to incrementally remove irrelevant aspects of a conceptual description and to automatically generate rules. We believe that these processes are applied to behaviors as well, for

³ It violates the *no structure in function* principle (de Kleer & Brown, 1984).

⁴ We think it is unlikely that this rule application method suffices for all cognitive processes that use rules; encoding seems to require an additional mechanism, for example. However, we wish to point out that the dividing line between rule-based processes and similarity-based processes may not be as solid as some might suppose.

the construction of prototypical behaviors (Forbus & Gentner, 1986) and situated rules.

Scaleable similarity-based retrieval. MAC/FAC (Forbus, Gentner, & Law 1995) models similarity-based retrieval. The MAC stage first uses a simple, non-structural matcher to filter out a few promising candidates from a (potentially immense) memory of structured descriptions. The FAC stage then evaluates these candidates more carefully, using a structural matcher (SME). The MAC stage lends itself to implementation in parallel (including connectionist) hardware, and has been tested with a variety of representations.

Analogical theory construction. PHINEAS (Falkenhainer 1987, 1990) demonstrated that structure-mapping processes could model several aspects of scientific theory construction, including matching behaviors and constructing qualitative domain theories by elaboration of candidate inferences. We believe that PHINEAS provides a good working model of some ways analogy is used in theory construction. Moreover, we conjecture that the same processes can be applied to modeling aspects of learning by instruction (where the teacher provides the starting correspondences) and within-domain analogical learning (where overall similarity guides the initial behavior match, and situation-specific explanations are incrementally generalized by a SEQL-like process).

Making Predictions

Let us see how these pieces might combine to solve a prediction task. Let the input be a (partial) description of a physical situation. An augmented version of *generate and test* could be used to make predictions as follows:

1. Retrieve similar behaviors (using MAC/FAC). The candidate inferences from mapping these remembered behaviors onto the observed behavior provide additional expectations about the current situation, and hypotheses about the states to follow, based on what happened in the remembered behavior. The state transitions hypothesized in the candidate inferences form the initial set of predictions.
2. If qualitative simulation rules or procedures are available for generating new behaviors (either by association with this type of task or because they are retrieved by MAC/FAC along with the behaviors used in the previous step), use them to expand the set of predictions.
3. If qualitative simulation rules or procedures are available for evaluating the consistency of possible transitions (from the same sources as the previous step), use them to filter the set of predictions.
4. If there are multiple predictions remaining, estimate their relative likelihood. Return the best, or several, if others are close to the best.

The first step provides quick recognition of familiar behaviors. If the overlap with the current situation is high and the behavior predicted unique, processing may stop at this point, depending on task demands⁵. The second step augments this recognition by domain-specific or first-principles consequences. The third step provides an opportunity for applying exceptions and caveats ("if it were overflowing, you would see coffee coming down the outside of the cup" and "strong acid dissolves coffee cups"). In the fourth step, we suspect that a variety of methods are used to estimate relative likelihood, ranging from domain-specific knowledge ("filling a wax paper cup with hot coffee usually causes it to leak") to estimates of relative frequency based on accessibility in memory ("I've never seen a ceramic coffee cup shatter when it was filled").

Psychological implications

Qualitative reasoning is not an island; it should use the same mental processes used in other aspects of cognition. Consequently, properties of analogical processing that have been found in other areas of cognition should appear in reasoning about mental models as well. We focus on three predictions next.

Distribution of reliance on memory with expertise

We conjecture that the use of memories in predictions with experience may vary as a U-shaped curve. That is, when little is known, memory use dominates, because comparison with previously observed behaviors encoded in perceptual terms are all that is available. As more is known, memory use may drop in favor of more abstract representations, such as situated rules. This may be especially likely in domains where the learner needs to articulate their models, e.g. situations where they are collaborating with others. It may be the case that as the domain becomes very familiar, memory use increases again, because the learner has experienced a large number of samples from the distribution of situations that occur. The theory-laden vocabulary learned by this stage may also greatly increase the frequency of relevant reminders (see below).

Differences in novice/expert retrieval patterns

The usual pattern in similarity-based retrieval (Gentner, Rattermann, & Forbus, 1993) is that retrieval is heavily based on surface properties (i.e., information about appearance and attributes of participating objects) rather than relational properties (i.e., causal arguments or abstract principles). In experts, however, the frequency of relational reminders in-

⁵ Consider deciding where to put a cup of coffee down on an uncluttered dining table versus deciding where to put it down in a very cluttered office.

creases (Novick, 1988). A possible explanation for this phenomena is that an expert's ability to encode phenomena in theory-laden terms provides additional overlapping vocabulary that helps the MAC stage find appropriate matches. For example, in solving physics problems, it has been observed that experts sort problems based on similarity in underlying principle, while novices sort problems based on similarity in the kinds of objects involved (Chi, Feltovich, & Glaser, 1981). Additional support for this explanation is provided by results suggesting that inducing subjects to encode materials more deeply increases the proportion of relational reminders (Faries & Reiser, 1988). The same phenomena should be observable in teaching people to make predictions in novel domains.

Factors that should promote expertise

Research on the role of comparison in development suggests two ways to speed up learning: *progressive alignment* and *inviting comparisons with relational language*.

Progressive alignment: (Gentner, Rattermann, Markman & Kotovsky, 1995) By exposing someone to a large number of very similar examples, their conservative learning mechanisms are more easily able to create the abstractions needed for transferable knowledge than if the same examples are interspersed with very different examples. Kotovsky and Gentner showed that experience with concrete similarity comparisons can improve children's ability to detect cross-dimensional similarity. Specifically, 4-year-olds' ability to perceive cross-dimensional matches (e.g., matching size symmetry with color symmetry) was markedly improved after experience with blocked trials of concrete similarity (blocks of size symmetry and blocks of color symmetry), as compared to control groups who received no training.

Inviting comparison with relational language. (Gentner & Rattermann, 1991) Giving a learner language for expressing a shared relational system can dramatically improve their ability to learn it via comparisons.

For example, Kotovsky and Gentner (in press) taught 4-year-olds labels for the relations of monotonic change ("more-and-more") and symmetry ("even"). During the training task, children learned (with feedback) to classify the stimuli as to whether they were "more-and-more" or "even." After this training, the children who were successful in the labeling task scored far better on a cross-dimensional version of the task than children without such training.

Applying these results to qualitative mental models yields three suggestions for how they might be learned more easily:

1. Show learners many situations varying in quantitative details but with identical qualitative behaviors

before showing them behaviors with a different qualitative structure. For example, someone learning about heat and temperature might first be exposed to a number of situations involving only heat flow before showing them a situation where heat flow is involved in phase changes, because in the latter the temperature of the object changing phase remains constant instead of increasing.

2. Name patterns of behavior (heating, cooling) first, and then move on to naming the physical mechanisms underlying them (heat flow, boiling).

3. Teach the compositional primitives of qualitative physics explicitly, to give learners a richer vocabulary for expressing their partial but growing knowledge.

Conclusions

Are qualitative mental models simulations or memories? Our answer is, some of both. No "pure" model provides a sufficient account for the runnability of human mental models. We propose instead a hybrid model, where similarity-based processes of comparison and abstraction provide the initial organization for knowledge of a domain, and broader principles of qualitative reasoning emerge from the accumulation and analysis of large numbers of examples, aided by the use of relational language as a focusing device and an invitation to comparison. As they emerge, these principles can be used for more rule-directed reasoning, but this augments, not replaces, analogical reasoning. We believe such a hybrid model is necessary to capture the flexibility of human common sense reasoning about the physical world over a broad range of states of knowledge. Currently we are exploring this model further, using a combination of psychological experimentation and computer simulation.

We would like to close with two points. First, we believe that qualitative simulation should not be identified only with reasoning from first principles using generic domain theories. The psychological intuitions that originally gave rise to the notion of qualitative simulation might be better served by making its defining characteristic be the use of qualitative representations to simulate, even if the predicted behavior is generated via analogical reasoning or with domain-specific rules. Second, we believe that qualitative physics has much to offer cognitive psychology. The vocabulary of qualitative physics (e.g., processes, influences, etc.) seems well-suited for expressing human beliefs about physical phenomena. We believe that an account combining these representational resources with analogical processing could provide a deeper understanding of human physical reasoning. Qualitative physics is already proving its worth in real-world applications. It can also contribute what might be a key piece in solving the puzzle of human cognition.

Acknowledgments

This paper has benefited from discussions with John Everett, Ron Ferguson, Tom Mostek, and Jeff Usher. This research is supported by the Cognitive Science Division of the Office of Naval Research, Contract N00014-92-J-1098.

References

- Bobrow, D., Falkenhainer, B., Farquhar, A., Fikes, R., Iwasaki, Y., Forbus, K., Gruber, T., and Kuipers, B. 1996. A compositional modeling language. *Proceedings of QR96*. AAAI Technical Report WS-96-01. AAAI Press.
- Bredeweg, B., and 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.
- Brown, J. S., Collins, A., and Duguid, P. 1989. Situated cognition and the culture of learning. *Educational Researcher*, 41:32-42.
- Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121-152.
- de Coste, D. 1994. Goal-Directed Qualitative Reasoning with Partial States. (Tech. Rep. No. 57). Evanston: Northwestern University, Institute for the Learning Sciences.
- de Kleer, J. & Brown, J.S. A qualitative physics based on confluences. *Artificial Intelligence*, 24: 7-83, 1984.
- Falkenhainer, B. (1987). An examination of the third stage in the analogy process: Verification-based analogical learning. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, August, 1987.
- Falkenhainer, B. (1990). A Unified Approach to Explanation and Theory Formation. Shrager & Langley (Eds.), *Computational Models of Scientific Discovery and Theory Formation*. San Mateo, CA: Morgan Kaufmann Publishers
- Falkenhainer, B., & Forbus, K. (1991). Compositional Modeling: Finding the Right Model for the Job. *Artificial Intelligence*, 51, 95-143.
- Falkenhainer, B., Forbus, K., Gentner, D. (1989). The Structure-Mapping Engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- Faries, J. M., & Reiser, B. J. (1988). Access and use of previous solutions in a problem solving situation. *Proceedings of the Tenth Annual Meeting of the Cognitive Science Society* (pp. 433-439). Montreal; Hillsdale, NJ: Erlbaum.
- Forbus, K. "Qualitative reasoning about space and motion" in Gentner, D. and Stevens, A. (Eds.), *Mental Models*, LEA associates, inc., 1983
- Forbus, K. Qualitative Process theory. *Artificial Intelligence*, 24, 1984.
- Forbus, K., Ferguson, R. and Gentner, D. 1994. Incremental structure-mapping. *Proceedings of the Cognitive Science Society*, August.
- Forbus, K. and Gentner, D. "Learning Physical Domains: Towards a Theoretical Framework" in Michalski, R., Carbonell, J. and Mitchell, T. *Machine Learning: An Artificial Intelligence Approach*, Volume 2, Tioga press, 1986.
- Forbus, K., Gentner, D., Everett, J. and Wu, M. 1997. Towards a computational model of evaluating and using analogical inferences. *To appear in Proceedings of the Cognitive Science Society*, August.
- 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., Nielsen, P. and Faltings, B. "Qualitative Spatial Reasoning: The CLOCK Project", *Artificial Intelligence*, 51 (1-3), October, 1991.
- Funt, B. 1980. Problem solving with diagrammatic representations. *Artificial Intelligence* 13 pp 201-230.
- Gambardella, L., Gardin, F., and Meltzer, B. 1988. Analogical representation in modeling naïve physics. *Proceedings of the 2nd International Workshop on Qualitative Physics*, Paris, France.
- Gentner, D. 1983. Structure-mapping: a theoretical framework for analogy. *Cognitive Science*, 23, 155-170.
- Gentner, D., & Markman, A. B. 1997. Structural alignment in analogy and similarity. *American Psychologist*
- Gentner, D. & Rattermann, M.J. 1991. Language and the career of similarity. In S.A. Gelman & J.P. Byrnes (Eds.), *Perspectives on language and thought: Interrelations in development*, Chap. 7 (pp 225-275). London: Cambridge University Press.

- Gentner, D., Rattermann, M.J., & Forbus, K.D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology* 25, 524-575.
- Gentner, D. and Stevens, A. (Eds.) 1983. *Mental Models*. LEA Associates
- Glasgow, J. 1992. Computational Imagery. *Cognitive Science* 16(3) 355-394.
- Guerrin, F. (1995). Dualistic algebra for qualitative analysis. In *Proceedings of 9th International Workshop on Qualitative Reasoning*, Amsterdam, Holland.
- Hegarty, M. 1992. Mental animation: Inferring motion from static diagrams of mechanical systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18:1084-1102.
- Hinton, G. 1979. Some demonstrations of the effects of structural descriptions in mental imagery. *Cognitive Science* 3:231-250.
- Johnson-Laird, P. 1983. *Mental Models*. Harvard University Press.
- Kosslyn, S. M. *Image and Mind*. Cambridge, Mass: Harvard University Press.
- Kosslyn, S. M. *Image and Brain*. Cambridge, Mass: MIT Press.
- Kotovskiy, L. & Gentner, D. (in press) Comparison and categorization in the development of relational similarity. *Child Development*.
- Kuipers, B. "Qualitative Simulation", *Artificial Intelligence*, 29, September, 1986.
- Kuipers, B. 1994. *Qualitative Reasoning: Modeling and simulation with incomplete knowledge*. Cambridge, Mass.: MIT Press.
- Mavrouniotis, M., & Stephanopoulos, G. (1988). Formal order-of-magnitude reasoning in process engineering. *Computers and Chemical Engineering*, 12 (9/10). 867-881.
- Norman, D. 1988. *The Psychology of Everyday Things*. Basic Books.
- Novick, L. R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 510-520.
- Price, C., Pugh, D. Wilson, M. and Snooke, N. The FLAME system: Automating electrical failure modes and effects analysis (FMEA). Proc. Annual Reliability and Maintainability Symposium, pp 90-95, IEEE 1995.
- Raiman, O. (1991). Order of magnitude reasoning. *Artificial Intelligence*, 51. 11-38.
- Rickel J., & Porter, B. (1994). Automated Modeling for Answering Prediction Questions: Selecting the Time Scale and System Boundary. *Proceedings of AAAI-94*. 1191-1198.
- Schwartz, D. 1996. Analog imagery in mental model reasoning: Depictive models. *Cognitive Psychology* 30, 154-219.
- Shimomura, Y., Tanigawa, S., Umeda, Y., & Tomiyama, T. (1995). Development of Self-Maintenance Photocopiers. *Proceedings of IAAI-95*. 171-180.
- 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.
- Spiro, R.J., Feltovich, P.J., Coulson, R.L., & Anderson, D.K. 1989. Multiple analogies for complex concepts: Antidotes for analogy-induced misconception in advanced knowledge acquisition. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp 498-531). Cambridge, MA: Cambridge University Press.
- White, B. & Frederiksen, J. 1990. Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence*, 42, 99-157.