

Dark Knowledge in Qualitative Reasoning: A Call to Arms

Kenneth D. Forbus¹, Dedre Gentner²

¹Qualitative Reasoning Group, Northwestern University

2133 Sheridan Road, Evanston, IL, 60208, USA

²Psychology Department, Northwestern University

2029 Sheridan Road, Evanston, IL, 60208, USA

{forbus, gentner}@northwestern.edu

Abstract

While people do qualitative reasoning, there is ample evidence that they do not always do it well. Two current crises, human-induced climate change and the financial meltdown, can be traced in part to faulty mental models. The QR community has formalisms that can potentially help with public education about such problems, but so far we have not been very successful in doing so. We claim that part of the reason is that current QR accounts do not adequately incorporate experiential knowledge. We argue that it is important to find better ways to improve public qualitative reasoning abilities, in part by helping people enlist their experience-based models via analogy.

Introduction

It is a truism that democracy works best with informed citizens. Alas, there is ample evidence that people do not have adequate mental models for many relevant areas. Consider two current crises: Human-induced climate change and the financial meltdown. In both cases, to be sure, there are people and organizations that are deliberately trying to obscure what is happening for their own reasons. But there is ample evidence that a fundamental failure of education has taken place. One key problem in understanding climate change is what Cronin *et al* (2008) call *stock-flow failure*. In system dynamics terms, a stock is something that accumulates, i.e., something that would be modeled in QP theory (Forbus, 1984) as a directly influenced parameter. A system with only flows has no accumulation, and its outputs are basically a function of its inputs. Surprisingly, when people are given graphs indicating the inputs to a system with accumulation, they often ignore the accumulation, and sketch the output as if it were simply a function of the inputs. This failure occurs even in highly educated people with technical backgrounds. Cronin *et al* further showed through a series of experiments that this could not be explained by problems in interpreting graphs, misunderstanding of context, lack of motivation, or lack of cognitive capacity. It is, quite simply, a failure of mental models reasoning.

What should be striking for our community is that we have what may potentially be some of the best ideas for helping people overcome these problems. It is difficult to teach ideas without accessible formalisms. The formalisms of QR, which factor out traditional mathematics and make causality explicit, could be of great value in education. But we have not been very successful in spreading these ideas more broadly.

This paper argues that to change this situation, we need to expand our models to be more psychologically oriented. Physicists postulated dark matter in order to explain the missing mass in their observations. By analogy, “dark knowledge” is the knowledge for which we lack elegant formalisms in QR, but which nevertheless is a major factor in human qualitative reasoning. Dark knowledge is concrete knowledge: specific facts and cases, derived from first-hand experience or via culture, that are remembered and used for many daily reasoning tasks via analogy. In terms of its size, we believe it far outweighs fully general first-principles knowledge, providing the “missing mass” that holds our conceptual universe together.

Understanding how people reason with dark knowledge is important for QR to reach its full potential. Moreover, we believe that understanding how human qualitative reasoning works is crucial for finding better ways to teach people to reason well about the complex problems we all face. This paper also argues that harnessing experiential knowledge through analogy is one potential way to transform education, making it better able to prepare people for the challenges ahead.

We begin by looking at the broad picture that QR and psychology paint of qualitative knowledge and reasoning. Next we look at how to improve human mental models, using a favored explanatory device for climate change, and one of QR’s favorite examples, the humble bathtub (Kuipers, 1994), to highlight how the use of analogy in explanations can be improved. After that, we discuss some ideas for tools and techniques for both understanding experiential knowledge better, and for improving education. We end with a call to arms.

Human mental models: The big picture

The study of qualitative reasoning was originally motivated by observing human reasoning: People who do not know differential equations reason about many physical phenomena perfectly well, and even scientists and engineers rely on simpler, qualitative models when framing problems and interpreting data. The “standard model” in QR explains this in terms of general, broad-coverage domain theories, expressed exclusively in first-principles terms. Given a particular scenario or problem, these general concepts are applied via instantiation to create a scenario model that can then be reasoned with.

This standard model has been remarkably successful in building a variety of useful systems. However, we believe it has strong limitations as a psychological account of human qualitative reasoning. We have proposed that much of human knowledge about the physical world is concrete (Forbus & Gentner 1997). In some sense it’s obvious: people have episodic memories, whereas most of today’s QR models do not. But there are reasons to believe that the use of experiential knowledge has profound consequences for human qualitative reasoning.

When someone starts learning about a novel phenomenon, they accumulate experiences. Even concrete experiences can be used directly in very similar situations, via within-domain analogies. We take experience quite broadly here: We include cultural influences such as language and education, as well as first-hand interaction with the world. While many of us have read about carbon sequestration and credit default swaps, for example, few of us have actually experienced these processes first-hand.

As experiences accumulate, they are used to construct generalizations, at first prototypical behaviors (*protohistories*, in Forbus & Gentner (1986)) and later causal fragments that can be turned into model fragments (the *causal corpus* in Forbus & Gentner (1986)). These generalizations are one source of misconceptions. Importantly, earlier forms of knowledge are added to, but not replaced by, later, more refined models. Once someone learns differential equations, for example, they still use simpler models, learned earlier, to throw balls, estimate stopping distances while driving, and other tasks where differential equations are in principle relevant.

Our hypothesis is that much of the knowledge people use in qualitative reasoning is concrete, at the level of protohistories and causal corpus. To be sure, we believe that something like first-principles domain theories are learned, either via analogical generalization or via direct instruction. In experts they are especially rich, including a tight integration with mathematical models. But even experts rely on experience-based models in their professional reasoning. For example, analogy seems to play an important role in model formulation (Falkenhainer, 1992; Klenk *et al* 2005). In non-experts, or even in experts, knowledge in many domains can be thought of as “pastiche models” (Collins & Gentner, 1987) or “in pieces” (diSessa, 1993), i.e., local, context-specific models.

Ideally, knowledge learned in school becomes tightly integrated with knowledge learned from experience, reorganizing it in ways that make correct reasoning more likely. Unfortunately, there is ample evidence that this integration is difficult, often leading to accumulation of multiple conflicting models. For example, Clement (1982) and McCloskey (1983) both showed that even students who did well in physics classes often continue to have and use incorrect qualitative models of force and motion. These misconceptions are uncorrelated with mathematical knowledge, and even honors students are susceptible to them (Halloun & Hestenes 1985). New misconceptions can arise during instruction as well (Spiro *et al* 1989; Vosniadou, 1994). Moreover, students sometimes actively work to protect and maintain their misconceptions, erecting “mental shields” when they are threatened by new information (Feltovich *et al* 2001).

Let us reexamine the stock-flow failure identified by Cronin *et al* (2008) in light of this model. They argue that people use a *correlation heuristic* in reasoning about systems with multiple continuous inputs and outputs. That is, when given the task of controlling a system which accumulates something, they tend to believe that the shape of the output should look something like the shape of the input, but delayed in time. This is the sort of heuristic that could very easily be derived from everyday experience, where the preponderance of input/output pairs we see are more often correlated in their behavior than not. If we turn the faucet in the sink or bathtub higher, water comes out faster, perhaps after some delay. The same thing happens when we turn on the tap on a garden hose. If this heuristic works in many situations, it is natural to apply it to new problems.

How can we improve mental models?

How can we improve people’s mental models? Simply handing them a modeling language, even in student-friendly terms (e.g. Betty’s Brain, Biswas *et al* 2001; VModel, Forbus *et al* 2004) is not enough. Showing them qualitative simulations (e.g., Bredeweg *et al* 2008) is not enough. These both are good starts, but unless we work on ways to integrate what they learn from these experiences into their prior knowledge, such interventions will not have as much impact as desired.

We believe that analogy is an excellent mechanism for integrating knowledge. Understanding the connections between experiences and/or models requires comparing them and understanding “what goes with what”, which is exactly what the structural alignment process at the heart of analogy does (Gentner, 1983). Further evidence indicating that analogy can be used to rapidly learn mental models (Kotovsky & Gentner, 1996; Gentner *et al* 2009) Combining the conceptual clarity of qualitative representations with the integrative power of analogy is, we suspect, exactly what we need to create new ways to help people reason better about complex situations. Showing how to think formally and qualitatively about

systems that someone has experienced first-hand provides a solid base domain that can then be projected by analogy to other target domains that need to be understood. Leveraging everyday experience provides solidity to conclusions that might not otherwise be plausible.

Research in psychology and learning sciences provides some insights for the effective use of analogy. For example, it is important to have learners work through correspondences in detail, so that they get the most out of the analogy (Kurtz *et al* 2001). Ensuring that the base domain is well-understood, and learners are focused on the relevant aspects of it, helps them apprehend the analogy (Richland *et al* 2007). We illustrate via an example next.

The bathtub analogy for climate change

Understanding climate change has proven to be quite difficult. Part of the problem is how counter-intuitive it is: For most of human history, people were at the mercy of weather and climate, and our impact seemed extremely small compared to the vastness of the planet. But as we grew in number and the planet did not, this changed. Now the modeling assumption of endless resources is clearly not accurate. One analogy that has been used to communicate the problem (e.g. Sterman 2008) is a favorite QR example, the humble bathtub.

Bathtubs have a faucet (or faucets) which can be opened to let water in, a drain which can be opened to let water out, and some capacity for holding water plus one or more people. Overflowing is something to be avoided. Our experience teaches us that for some level of water, it is likely that when we sit down in the tub it will overflow. In this analogy, the atmosphere is like the volume of the bathtub. The accumulation of carbon in the atmosphere is like the accumulation of water in the bathtub. Just as there is a level at which overflows are likely in the bathtub, there is some level at which accumulated carbon causes problems on a massive scale (countries going underwater, starvation, etc.).

In explaining this analogy so far, we have been very explicit about what aspects of the base domain should be considered, so explicit as to cross the line into belaboring the obvious for the already-informed. This degree of elaboration is useful to provide a solid foundation for extending an analogy into new areas, or using it to help understand new ways of reasoning. Having students work through the correspondences explicitly and in detail, by constructing a table for instance, helps ground the mapping.

This analogical model provides considerable value in reasoning. If the inflow is larger than the outflow, then the level will be rising. This is what is happening in the atmosphere, with CO₂, methane, and other greenhouse gasses being produced faster than natural processes can absorb them, and hence they accumulate in the atmosphere. Opening the tap wider in the bathtub will cause the level to rise faster, and increasing carbon emissions will lead to disaster more quickly. In public policy terms, a “conservative” strategy often proposed is to keep carbon

emissions at their current level. But, mapping this strategy to the bathtub, one can easily see that this is not enough: The level will continue to rise inexorably to overflow, unless emissions are reduced below their current levels.

A good analogy provides a framework that can be expanded to incorporate additional ideas. For example, suppose we cannot or will not turn down the faucet. The only way to prevent an overflow is to increase the rate of draining – with buckets, if need be. In the case of the atmosphere, planting new forests is one way to improve its “drainage”. Unfortunately, a recent result about the oceans absorbing less carbon due to increases in atmospheric temperature¹ can be understood as one of the “drains” becoming less effective, and thus leading to a higher rate of carbon accumulation – a potentially nasty positive feedback cycle.

Modeling bathtubs, and other everyday examples, is a common practice in QR because it allows us to compare formalisms more easily. It is also a useful exercise for someone learning a new modeling language because it helps integrate the new language into their experiences. It is important to walk through everyday behaviors, and show how they can indeed be derived from the consequences of the primitives. For instance, the relative rate of the inflow and outflow determines whether the amount of water, and hence the level, is increasing. The idea that one can get a stable balance between inflow and outflow for a range of levels can also be examined, although this will take more work since people are less likely to be familiar with this notion. To see how important elaborating the everyday example is, consider this: In some of the experiments exploring the stock-flow failure, the simulated system being controlled was a bathtub! When entering a technical problem, people often check their intuitions at the door. Tightly coupling abstract models and everyday experience seems central to the challenge we face.

Promoting transfer

How can we help people apply new ideas when they are potentially relevant? Research on analogy in instruction suggests that having learners compare cases can double the odds of them applying concepts to new situations where they are relevant (Gentner *et al* 2003). Re-using the bathtub as an analog to credit card debt provides an example.

Bathtub	Credit Card Usage
Faucet setting	Monthly charges
Drain setting	Monthly payment
Level of water	Amount of debt
???	Interest rate

This table of correspondences helps us understand that we are missing something in the analogy: What is the

¹ <http://www.guardian.co.uk/environment/2009/jan/12/sea-co2-climate-japan-environment>

bathtub equivalent of the interest rate on a credit balance? This is like a second faucet, whose setting is determined by the level of water and the interest rate. So even if there are no new monthly charges, debt will continue to accumulate, thanks to this second faucet. Again, this may seem obvious, but it is interesting just how many people in the US economy behave as if they do not believe this is true.

Mental models and the financial crisis

Untangling the causes of the current financial crisis is an ongoing process, being undertaken from a variety of perspectives by a number of disciplines. Consequently, the evidence here is less well worked out than in the case of climate change. However, even at this stage of understanding, some reasonable conjectures can be made. One factor appears to be the seduction of mathematical models, especially embedded in software, over historical experience. Markets go down as well as up. The history of economic bubbles provides ample evidence that people tend to ignore this fact (Mackay, 1841). Coming up with an accurate and clear model of this debacle is itself a tough qualitative modeling challenge. For example, one of the factors that has made the current crisis so widespread is the dangerous process of “repackaging” mortgages as if they were securities. If qualitative models of causal factors affecting risk were included, and propagated through the multiple levels of repackaging, they might have helped alert investors to the potential dangers. Such models will require reasoning about distributions – if the economic climate becomes tougher, business will lay off employees. If many people are laid off, then they will be unable to pay their mortgages², leading to the collapse of these “securities.” This in turn makes the economic climate tougher still, by drying up credit. Being able to systematically examine worst-case, as well as best-case, possible outcomes might help mitigate the “herd thinking” that underlies bubbles.

What is to be done?

We believe that the QR community has unique contributions to make in helping to improve public education on climate change, financial problems, and other issues raised by our more complex and more tightly interconnected world. We see a three-pronged approach as necessary: (1) more research on the nature of human mental models, including experiential knowledge, (2) more research on how to improve human learning and reasoning, and (3) construction of tools that help people reason and learn, based on the best available results from cognitive science (including learning sciences). We consider each in turn.

² As of 3/5/09, 48% of Americans with subprime mortgages are behind on their payments or are in foreclosure. Source: http://www.newsvine.com/_news/2009/03/05/2508945-mortgage-woes-break-records-again-in-4q

Understanding human mental models

In the 1980s and early 1990s, much of the energy in the QR research community was spent on developing formalisms for qualitative dynamics. While the accounts developed have been shown to be robust, by being used in a wide variety of problems and domains, the climate and financial crises illustrate that either (a) these formalisms are not being used by people or (b) there are other representations and processes being used in human reasoning as well. The evidence against (a) mostly comes from protocol analyses, and more research establishing that people do in fact use ideas like qualitative proportionality to organize causal models is needed. The evidence for (b) is strong, e.g. the misconception literature in science education. We believe that the nature of experience-based knowledge must be better understood, and that no account of qualitative reasoning and its place in common sense will be complete without it.

Another reason for strengthening qualitative reasoning skills is to overcome the blind acceptance of the authority of mathematical models. In the financial crisis particularly, executives relied on models produced by their “quants” without fully understanding their implications. Better articulation of the underlying assumptions and causal factors assumed might have led to more caution.

Psychologists have an easier time exploring experience-based knowledge because they can study systems that have plenty of it (i.e., people). For computational modeling, the situation is more complicated. Most QR systems are either fed their knowledge by hand, or are processing information from a specific set of numerical sensors. Hand-feeding systems descriptions expressed in their internal representations does not scale very far. Exploring the role of experience in qualitative reasoning requires finding reasonable approximations to the representations that people build up by interacting with the world. Importantly, by “world” we mean both the physical world and the cultural world: Many physical phenomena are only experienced at best indirectly, with our models of them gleaned from our culture, via reading, lectures, and conversations.

As progress in vision and robotics continues, there will be platforms where experience can be directly gathered by interacting with the physical world. But we need not wait for such platforms, especially given the importance of cultural inputs in human learning. It is already possible to create systems that semi-automatically produce formal representations from simplified natural language (e.g. Kuehne & Forbus, 2004; Tomai & Forbus 2009) and sketches (e.g. Forbus *et al* 2008). These media are relatively easy to produce, and can be used to experiment with learning experience-based models (e.g., Friedman & Forbus 2008; Friedman & Forbus 2009).

Improving human learning and reasoning

The misconception literature in science education shows that helping people achieve accurate models of physical

phenomena is quite difficult. For many problems that become matters of public policy there are two additional sources of difficulty: (1) They are more complex, in terms of the number and variety of causal influences and (2) There can be vested interests actively attempting to sow confusion, to better achieve their own ends. When education becomes an adversarial game, it becomes much harder.

Here is an example: George Will, in the Washington Post on 2/15/09, wrote

“As global levels of sea ice declined last year, many experts said this was evidence of man-made global warming. Since September, however, the increase in sea ice has been the fastest change, either up or down, since 1979, when satellite record-keeping began. According to the University of Illinois' Arctic Climate Research Center, global sea ice levels now equal those of 1979.”

There are a number of misstatements here. The first is a misrepresentation of how data are evaluated. As Andrew Revkin, in a New York Times blog posting³, puts it,

“No single year marks a trend or holds evidence of long-term climate change.”

He quotes Jennifer Francis, from Rutgers, who responds to one of Will's assertions with an excellent qualitative explanation of why warming has contributed to the speed of ice recovery:

“At the end of summer each year, the sea ice refreezes and continues to do so until late spring. Thin ice and open water generate new ice faster than thick ice, as the heat from the ocean below is able to escape more easily to the atmosphere. In the autumns of 2007 and 2008, the rate of ice production was very large because there was so much open water and thin ice – the rapid growth is completely expected.”

Mr. Will's confusion is symptomatic of a major problem we have in our culture. When journalists and opinion-makers have trouble understanding scientific evidence and arguments, the effect of their confusion is multiplied by decreasing the clarity of public debate.

Building tools to support reasoning and learning

The QR community has already invested substantial effort into making tools that use qualitative modeling to help students learn and to help inform the public about the possible consequences of policy choices (e.g., Sallas & Bredeweg, 2001). There is certainly much more to be done in this area.

The importance of experiential knowledge in human qualitative reasoning suggests that we need to incorporate ways to exploit it into our tools. For example, our QCM system (Dehghani & Forbus, 2009) is a new qualitative

modeling tool aimed at cognitive scientists, to help them model data that they have collected. It deliberately allows users to create situation-specific descriptions of physical processes, rather than forcing them to first create and then instantiate a first-principles domain theory. The idea is that situation-specific models may be all that they need for particular investigations, and that even if their goal is to construct a robust, broadly-applicable first-principles domain theory for some area of human knowledge, building concrete, specific models is a better way to start. In other words, contemplating multiple specific models may be a better way to formulate general domain theories. We suspect that the same approach could be useful for students as well, given the success of Betty's Brain and VModel.

Another way to incorporate experiential knowledge in our learning and reasoning tools is to enable them to work with their users' analogies, and to supply their own. Explicitly helping people work through correspondences and seeing what analogical inferences follow, for example, could be a valuable service in a learning environment. A system could propose new analogies, drawing upon interesting examples it has formally represented as part of its world knowledge. People can often work through an analogy once it is proposed, but they find it much harder to retrieve distant (as opposed to close) analogs (Gentner, Rattermann, & Forbus 1993). Support software can potentially have an easier time retrieving distant analogs, since they have fewer distracting experiences, fewer distracting perceptual representations, and can encode experiences thoroughly.

A Call to Arms

We believe that the ideas and formalisms developed by the qualitative reasoning community can play an important role in public education. Democracies require informed citizens. In today's world, citizens are faced with the need to understand quite subtle arguments about very complex interlocking systems, and have to sift through both honestly conflicting evidence and special-interest induced hazes. We believe that the ability to do robust, sound qualitative reasoning is an important part of meeting this need. But to succeed, we must take into account experiential knowledge, the “dark knowledge” of QR, because it seems to play a central role in human mental models. Our models of reasoning and learning need to incorporate it, and our designs for educational systems and interventions need to take it into account.

Advances in natural interaction modalities (natural language, sketch understanding, vision, robotics) provide new tools by which we can accumulate in digital form knowledge about experience. We hope that this will facilitate research on the roles of experience in qualitative reasoning. And we hope that this, in turn, will help us develop a new generation of QR techniques and systems to help with these crucial matters of public education.

³ <http://dotearth.blogs.nytimes.com/2009/02/27/expers-big-flaw-in-wills-ice-assertions/>

Acknowledgements

We thank Morteza Dehghani, Scott Friedman, Andrew Lovett, and Jon Wetzel for helpful suggestions.

References

- Biswas, G., Schwartz, D., Bransford, J. & The Teachable Agents Group at Vanderbilt. (2001) Technology Support for Complex Problem Solving: From SAD Environments to AI. In K. Forbus and P. Feltovich (Eds.). *Smart Machines in Education*, 71-98. AAAI Press/MIT Press
- Bredeweg, B., Salles, P., Bouwer, A., Liem, J., Nuttle, T., Cioaca, E., Nakova, E., Noble, J., Caldas, A., Uzonov, Y., Varadinova, E., and Zitek, A. 2008. Towards a structured approach to building qualitative reasoning models and simulations. *Ecological Informatics*, 3(1), 1-12. Elsevier.
- Clement, J. 1982. Students' preconceptions in introductory mechanics. *American Journal of Physics* 50(1):66-71.
- Cronin, M., Gonzalez, C., & Serman, J. 2009. Why don't well-educated adults understand accumulation? A challenge to researchers, educators, and citizens. *Organizational Behavior and Human Decision Processes* 108, 116-130.
- Dehghani, M. and Forbus, K. 2009. QCM: A QP-based concept map system. *Proceedings of QR09*.
- diSessa, A. (1993). Toward an Epistemology of Physics. *Cognition and Instruction*, 10(2-3), 105-225.
- Falkenhainer, B. 1992. Modeling without amnesia: Making experience-sanctioned approximations. *Proceedings of QR92*, 44-55.
- Feltovich, P., Coulson, R., and Spiro, R. 2001. Learners' (Mis)understanding of important and difficult concepts. In Forbus, K., and Feltovich, P. (Eds.) *Smart Machines in Education*, 349-375, AAAI Press/MIT Press.
- Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85-168.
- Forbus, K., Carney, K., Sherin, B. and Ureel, L. 2004. VModel: A visual qualitative modeling environment for middle-school students. *Proceedings of the 16th Innovative Applications of Artificial Intelligence Conference*, San Jose.
- Forbus, K. and Gentner, D. (1986). Learning Physical Domains: Towards a Theoretical Framework. In Michalski, R., Carbonell, J. and Mitchell, T. (Eds.), *Machine Learning: An Artificial Intelligence Approach*.
- Forbus, K., Usher, J., Lovett, A., Lockwood, K., & Wetzel, J. (2008). CogSketch: Open-domain sketch understanding for cognitive science research and for education. *Proceedings of the Fifth Eurographics Workshop on Sketch-Based Interfaces and Modeling*. Annecy, France.
- Friedman, S. & Forbus, K. (2008). Learning Causal Models via Progressive Alignment & Qualitative Modeling: A Simulation. *Proceedings of CogSci 2008*.
- Friedman, S., Forbus, K., 2009. Learning Naïve Physics models and misconceptions. *Proceedings of CogSci09*.
- Gentner, D. 1983. Structure-Mapping: A theoretical framework for analogy, *Cognitive Science* 7(2):155-170.
- Gentner, D., Levine, S., Dhillon, S. & Poltermann, A., (in press) Using structural alignment to facilitate learning of spatial concepts in an informal setting. To appear in *Proceedings of the Second International Workshop on Analogy*, Sofia, Bulgaria, 2009
- Gentner, D., Loewenstein, J., Thompson, L. 2003. Learning and transfer: A general role for analogical encoding. *Journal of Educational Psychology*, 95, 393-408.
- Gentner, D., Rattermann, M.J., & Forbus, K. 1993 The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology* 25, 524-575.
- Halloun, I. & Hestenes, D. 1985 Common-sense concepts about motion. *Am. J. Phys.* 53, 1056.
- Klenk, M., Forbus, K., Tomai, E., Kim, H., and Kyckelhahn, B. (2005). Solving Everyday Physical Reasoning Problems by Analogy using Sketches. *Proceedings of 20th National Conference on Artificial Intelligence (AAAI-05)*, Pittsburgh, PA.
- Kotovsky, L., & Gentner, D. 1996. Comparison and categorization in the development of relational similarity. *Child Development*, 67, 2797-2822.
- Kuehne, S. and Forbus, K. (2004). Capturing QP-relevant information from natural language text. *Proceedings of the 18th International Qualitative Reasoning Workshop*, Evanston, Illinois, August.
- Kuipers, B. 1994. *Qualitative Reasoning: Modeling and simulation with incomplete knowledge*. Cambridge, Mass.: MIT Press.
- Kurtz, K., Miao, C., & Gentner, D. 2001. Learning by analogical bootstrapping. *Journal of the Learning Sciences*, 10(4):417-446.
- Mackay, C. 1841. *Extraordinary Popular Delusions & the Madness of Crowds*. Three Rivers Press.
- McCloskey, M. 1983. Naïve theories of motion. In Gentner, D. and Stevens, A., (Eds) *Mental Models*. 299-324. Erlbaum.
- Richland, L., Zur, O., and Holyoak, K. 2007. Cognitive supports for analogies in the mathematics classroom. *Science* 316:1128-1129, May 25th.
- Salles, P. & Bredeweg, B. 2001. Constructing Progressive Learning Routes through Qualitative Simulation Models in Ecology. *Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future*, J.D. Moore, G. Luckhardt Redfield, and J.L. Johnson (eds), pages 595-597, IOS-Press/Ohmsha, Japan, Osaka.
- Spiro, R., Feltovich, P., Coulson, R., & Anderson, D. 1989. Multiple analogies for complex concepts: Antidotes for analogy-induced misconception in advanced knowledge acquisition. In *Similarity and Analogical Reasoning*, eds. S. Vosniadou & A. Ortony, 498-531. Cambridge University Press.

Sterman, J. 2008. Risk Communication on Climate: Mental Models and Mass Balance. *Science*, 322(532-533), Oct 24.

Tomai, E. and Forbus, K. (2009). EA NLU: Practical Language Understanding for Cognitive Modeling. *Proceedings of the 22nd International Florida Artificial Intelligence Research Society Conference*. Sanibel Island, Florida.

Vosniadou, S. (1994). Capturing and modeling the process of conceptual change. *Learning and Instruction*, 4, 45-69.