Appears in the Proceedings of the eighth annual meeting of the Cognitive Science Society, Amherst, MA, August 1986.

Causal reasoning about quantities

Kenneth D. Forbus
Qualitative Reasoning Group
University of Illinois

Dedre Gentner
Psychology Department
University of Illinois

Abstract

Causality plays an important role in human thinking. Yet we are far from having a complete account of causal reasoning. This paper presents an analysis of causal reasoning about changes in quantities. We abstract from AI theories of qualitative physics three dimensions along which causal reasoning about quantities may be decomposed. We then use this framework to make some psychological predictions.

1. Introduction

People have a deep intuition that causality is a central and cohesive aspect of human mental life. Consequently, the problem of causality has long occupied philosophers and scientists. But the search for a unified theory that can explain human causal reasoning, much as theories of grammar explain syntactic processing, has so far been unsuccessful. These failures have led Hayes (1985) and others to conclude there is no deep theory of causality. Instead, causal reasoning may be simply a family of inferences whose properties will vary according to the content of the argument. This paper analyzes one kind of causal argument, causal reasoning about changes in quantities, to provide an account of the relevant issues and draw some implications for psychology.

In this paper we abstract from the AI qualitative physics literature three factors involved in causal reasoning: (1) whether there is an *explicit mechanism* or not; (2) whether the connective relations between quantities have a *direction of causation* built in, and (3) which type of *measurement scenario* is involved. We begin by laying out these factors as they apply in current work in qualitative physics. We analyze the relationships among these factors and discuss the current literature in light of these distinctions. Then we discuss their implications for psychology. We advance some conjectures and consider suggestive evidence from protocols.

2. Models of causality in changing quantities

One of the central concerns of qualitative physics (e.g., Bobrow, 1985) is describing how continuous physical properties change over time. Informally, it seems that people treat many of these deductions as causal:

"Pouring more milk into the glass will cause the level to go up."

"Turning on the stove increases the temperature of the burner, which causes heat to flow, which eventually causes the water on the stove to boil."

Workers in qualitative physics have made a number of proposals about how to model these conclusions. In order to compare these proposals, we isolate three important factors. Any theory of causation involving quantities must make some choice within each of these dimensions. Thus we have a basis for comparing theories and organizing psychological predictions. These three factors are:

1. *Explicit/Implicit Mechanisms:* Does the theory include an explicit notion of mechanism, in addition to objects, that is the root cause of all changes?
2. *Directed/Nondirected Connectives:* Are the relationships between quantities expressed by functional dependencies whose directionality is taken to express the direction of causation, or by nondirected constraint equations?
3. *Measurement Scenarios:* What sense of change is being discussed? Is it one change in a sequence, the difference between initial and final states, the difference between alternate possible worlds, or something which is occurring continuously?

We examine each aspect in turn, noting the issues involved and how the current systems of qualitative physics deal with them.

2.1. Mechanism: Explicit versus Implicit

The issue of mechanism in physics arises in a subtle way. Traditional physics expresses many ideas informally since it can draw on our commonsense view of the world. Qualitative physics provides ways to formalize some of these ideas. One such aspect is expressing when different equations are valid. For example, the equations that describe the relationship between volume and temperature are different for a piece of ice, some water in a glass, and steam in a pressure cooker, even though the substance is the same in each case. There are several ways to formalize this knowledge, each involving different levels of ontological commitment.

Bare logical implication is one extreme alternative. Given the right predicates one can correctly specify when equations are applicable, but this alternative provides no organizational structure for physical knowledge. Thus most systems of qualitative physics provide some organizing mechanism, and we shall not consider this alternative further.[1]

The other extreme is to add mechanism, i.e. a special ontological class (or classes) for things that are to be the "agencies of causation". All changes are then stipulated to be directly or indirectly caused by some member of this class of mechanisms. We call these *explicit mechanism* accounts. The current explicit mechanism theories in qualitative physics use continuous or discrete *processes* as the class of mechanisms (Forbus, 1981, 1984; Simmons 1983; Weld 1984). Examples of processes include liquid flow, heat flow, and boiling. These processes operate on objects in various ways, causing changes to occur in them.

A middle position is that there is some mechanistic connection between parameters but there is no separate agency apart from the objects themselves. We call these *implicit mechanism* accounts. Current implicit mechanism accounts organize their laws around *devices* (de Kleer & Brown, 1984; Williams, 1984). Examples of devices include resistors, capacitors, and transistors.

---

[1] An apparent exception is the system of Kuipers (1984), but his goal is to produce a new qualitative mathematics compatible with any scheme for structuring the equations. Based on personal communication, we attribute to him the explicit mechanism position.

A system in the world is modeled by connecting together collections of device models into a network, and changes arise as a consequence of components interacting with other parts of the network.

At first glance explicit–mechanism theories might seem more complex since they posit extra entities. However, they can in fact simplify reasoning. Explicit–mechanism accounts facilitate making and using closed world assumptions (Collins et. al. 1975), which are necessary for beings with finite knowledge and computational resources. An example of a closed–world assumption in causal reasoning is "If nothing in the class of mechanisms I know about is causing a change, then the change cannot occur." If this assumption is violated, an explicit–mechanism account provides a possible way out, namely to postulate a new member of the class of mechanisms. Having a theory of mechanisms limits the search space when faced with contradictions.

2.2. Connectives: Directed or nondirected

The second aspect of causal reasoning about quantities concerns the relationship between the form of qualitative laws and their role in causal reasoning. The formal language of traditional physics is mathematics, typically differential equations. Clearly a qualitative physics must include some qualitative rendering of differential equations, and all of them do. But there are two choices for how equations are used to express causality:

1. *Directed connectives:* The qualitative equations are written as functional dependencies, where the direction of dependence is identified with the direction of causality, i.e.

   $$Q_0 = F(Q_1, Q_2 \ldots Q_n)$$

   will sanction the inference that a change in one of $Q_1, \ldots, Q_n$ can cause a change in Q_0, but not vice–versa.[2] For example, we might write Newton's second law as

   $$a = F / m$$

   to express that we can cause the acceleration to change by changing the force we apply or the object's mass, but not the other way around.

2. *Nondirected connectives:* Qualitative equations are written as constraint equations, and if there are n terms and n-1 of them are known, then the nth term can be calculated. Furthermore, no matter which quantities are involved, logical dependence can be interpreted as causation. For example, in electricity Ohm's law,

   $$V = I * R$$

   states that the voltage across a resistor equals the product of the current and the resistance. We can change the current to cause the voltage to change, and change the voltage to cause the current to change.

The difference between the two positions may be difficult to see at first, since functional dependency implies logical dependency and any constraint equation may be written as a function. The critical fact is that any constraint equation can be written as n different functions, where n is the number of variables in the equation. The directed connective position is that, while each of these different functions may be used in reasoning, only one of them will be distinguished as causal. With directed connectives the role of a qualitative law in causal reasoning is determined by its form, with nondirected connectives the role is determined by how it is used.

---

[2] This choice is not identical to the classical functional view of causation introduced by Mach (as described in Bunge (1979)), because not all functions are identified as causal.

In the current systems of qualitative physics, the choice of connectives has been more or less identified with the choice of explicit versus implicit mechanism. Explicit mechanism theories tend to use directed connectives, and implicit mechanism theories tend to use nondirected connectives. The exceptions are Kuipers (1984) and Williams (1984), who use both. Table 1 shows the 2 X 2 set of possibilities generated by crossing the Mechanism and Connectives dimensions.

We think the reason for this correlation is that explicit mechanism theories provide a notion of "independent" parameters, those directly affected by some mechanism. Effects then propagate outward from these distinguished parameters, like the level of water in a cup changing in response to pouring more water into it. The mechanism thus imposes the direction of causality on the system. Implicit mechanism theories do not identify independent parameters in advance (but see below), so it is hard to pre–judge the way a law will be used causally. Clearly there is no logical barrier to theories which inhabit any and all cells of this table, and later we describe combinations we believe may play signficant roles in human causal reasoning.

Each choice has problems. With nondirected connectives, causal reasoning requires an initial perturbation (such as increasing an input voltage to a circuit). However, people also have causal intuitions about situations even when they do not see the initial perturbation of the objects involved. For example, people are willing to say that the steam they see coming out of a kettle is caused by the boiling occurring inside it, even when they did not see the stove being turned on.

Furthermore, even if an initial perturbation is provided, it seems that some annotations about an equation's causal role are still necessary to avoid inappropriate causal inferences.

Table 1 – Possibilities for Mechanisms and Connectives
The choices along the mechanism and connectives dimensions have not been independent in systems of qualitative physics. The "*" indicates a system which predominately lies in that cell, but allows the other kind of connective as well.

Connectives	Mechanism: Explicit	Mechanism: Implicit
Directed	Forbus, Kuipers*, Simmons, Weld	
Nondirected		de Kleer & Brown, Williams*

Certain ways of using equations do not correspond to our intuitions about causation.[3] Returning to Ohm's law, we do not assume that increasing the current causes the resistance to change. Even theories which use nondirected connectives must, it seems, rely on some sort of annotation about direction. For example, in de Kleer and Brown's ENVISION program, we find in the description of their confluence heuristic (ibid, page 73),

> "In the specific case of the valve, the converse (where the area is changed) is impossible as the area is an input–only variable of the valve."

This concept of an "input–only" variable constitutes an annotation that violates the nondirectedness of their equations. Unfortunately, there is no theoretical guidance in their account as to which parameters should be so marked.

Directed connectives also have problems. They require the model builder to explicitly state which way causality works in all circumstances. Unfortunately, certain laws can be used causally in more than one direction. Ohm's law, once again, is a good example. If we are reasoning about a voltage source we want to make V be the independent variable, i.e., a change in voltage causes a change in current. If we are reasoning about a current source we want to make I be the independent variable, i.e., a change in current causes a change in voltage. It appears that in principle one can create models that use multiple directed connectives to capture these different causal interpretations, by explicitly specifying a context for each direction. However, writing nondirected constraint equations appears much simpler for these cases.

Some systems, notably those of Williams and Kuipers, attempt to circumvent these difficulties by allowing a mixture of directed and nondirected connectives. The obvious advantage is that the modeler is then free to choose whatever connective seems appropriate. Such freedom, however, can be dangerous. So far no mixed system has provided theoretical constraints on the choice of connective type, which means the choice must be made on an ad hoc basis. We believe such theoretical constraints probably exist, and could be generated by extending an explicit mechanism account.

2.3. Measurement Scenarios

The third aspect of qualitative causal reasoning concerns the sense in which a quantity is said to be changing. Consider again a kettle half–filled with water sitting on a stove. Suppose at some time T_0 we turn the stove on. At some later time T_1 the water begins to boil, and at time T_2 the water has completely boiled away. Most people would agree that from T_0 to T_1 there is a heat flow that causes the heat of the water and its temperature to rise, and that from T_1 to T_2 the boiling is causing the amount of water to decrease and the amount of steam in the room to increase. However, there are four different ways that one might discuss changes in quantities, even within this simple scenario:

1. *Incremental measurements.* We can think about what happens at T_0 when the flow begins. We could say,

   > "The temperature difference between the water and the burner causes heat to flow between them. This will then cause the heat of the water to rise, which will then cause the temperature of the water to rise"

---

[3] See Forbus (1984) for a detailed discussion.

The incremental scenario takes a sequential view of the property changes, demanding that one change occurs before another. In essence, this scenario extends the kind of causality we use for macroscopic discrete events (such as a row of dominos falling in succession after one is pushed over) into the realm of continuous changes. A prototypical example in the continuous realm is following a "piece of liquid" through a hydraulic system.

2. *Discrete measurements:* We can think about the difference between the world at T_0 and at T_2, without considering what happened in between. For example, we might note that there is now no water in the kettle, and the room we are in is more humid than when we started.

3. *Differential measurements:* We can think about what would happen if some property of the situation were different. For example, we might conclude that increasing the temperature of the stove would cause the steam generation rate to increase, and thus the water would boil away sooner. Essentially, we are comparing two possible worlds, related to each other by some change in property or occurrence.

4. *Continuous measurements:* We can think about what is happening during some particular kind of activity. For example, we can say during the interval between T_0 and T_1 that the increase in the heat of the water is causing its temperature to increase, even though both changes are occurring at the same time.

Each of these measurement scenarios has been used in qualitative physics. The incremental scenario was first introduced by de Kleer (1979), and is also used by Williams (1984). The discrete scenario has been used by Simmons (1983) and Weld (1984). The only specific proposal involving the differential scenario is *differential qualitative analysis* (see Forbus, 1984), but it is still relatively unexplored. The continuous scenario is used by most current systems of qualitative physics, including (de Kleer & Brown, 1984; Forbus, 1984; Kuipers, 1984; Williams, 1984).

The discrete, continuous, and differential views each have their distinctive role to play in reasoning about quantities. As argued in (Simmons, 1983; Weld, 1984), often the details of how some change occurs are unclear or irrelevant. A lower-precision discrete model which represents only end-states may best match the available information[4]. Conversely, the continuous view becomes essential when we are concerned with what is happening during a particular activity. The differential view provides information about how things would turn out differently if some change were made, and thus is useful in debugging.

The incremental scenario has considerable intuitive appeal. Unfortunately, so far this scenario has been problematic as a formal model. It requires a distinct notion of time, called *mythical time*. Unlike standard theories of time, mythical time is only partially ordered, and no real time passes between instants of mythical time. While some attempts to clarify the nature of mythical time and its relationship to normal time have been made (de Kleer & Brown, 1984; Williams, 1984) there is still no full formal account. Nevertheless, the incremental scenario is very important for psychological accounts of causality.

3. Psychological Implications

Table 2 summarizes the set of distinctions we have made. It can be seen that there are 16 theoretically possible cells, of which 5 contain AI qualitative physics theories. The "unified causal

---

[4] Domains for which the best models are discrete are outside the scope of this paper.

Table 2 – The space of causal theories about quantities

	Mechanism			
	Explict		Implict	
Measurement Scenario	**Connectives**		**Connectives**	
	Directed	Nondirected	Directed	Nondirected
Incremental				de Kleer & Brown, Williams*
Discrete	Simmons, Weld			
Differential	Forbus			
Continuous	Forbus, Kuipers*			de Kleer & Brown, Williams*

theory" approach to human causal reasoning would be to ask which cell is the one humans use. The questions raised by the "inference family" view of causation are more complex. The questions become:

1. Which cells do people typically use?

2. Are there characteristic patterns of use, such as novice–expert differences?

The remarks which follow are speculative, because these factors have not previously been fully isolated and subjected to systematic psychological investigation. Therefore what follows is a set of conjectures made in the hope of getting the empirical ball rolling. We begin with the first two factors, *explicit vs. implicit mechanisms* and *directed vs. nondirected connectives.*

These dimensions are highly correlated in theories of qualitative reasoning. Table 2 shows the concentration of AI theories in the two outer columns: explicit mechanisms with directed connectives, or implicit mechanisms with nondirected connectives. Here we ask how each of the four possible combinations of explict/implicit mechanisms and directed/non–directed connectives (i.e., the columns of Table 2) might be manifested in human reasoning about quantities. We suspect that (a) examples of all four combinations can be found in human reasoning; (b) the outer two columns, heavily explored in qualitative physics, do in fact represent common human causal arguments; but (c) the implicit mechanism/directed connective combination, unexplored by qualitative physics, also represents an important class of human causal reasoning.

We will illustrate the four classes with the familiar domain of car engines.

1. *Explicit mechanisms with directed connectives:* e.g.

> "Opening the throttle increases the flow rate of gas to the engine, which causes the engine to work faster."

Here the reasoner uses a set of processes to make causal inferences. This is the kind of reasoning modeled by Qualitative Process theory.

2. *Implicit mechanisms with directed connectives:* e.g.,

> "Driving faster causes fuel consumption to increase."

Instances of this class are sometimes instances of diSessa's *phenomenal primitives* or of what we call the Causal Corpus (diSessa, 1983; Forbus & Gentner, 1986).

3. *Explicit mechanisms with non-directed connectives:* e.g.,

> "Inside the engine, during the compression stroke the decrease in volume inside the cylinder causes the pressure to increase. During the expansion stroke, the increased pressure due to combustion pushes the cylinder down, causing the volume to increase."

4. *Implicit mechanisms with non-directed connectives:* e.g.,

> "The increased voltage at the input causes the current through resistor R_1 to rise. Since resistor R_2 is connected to resistor R_1, this increased current will cause the voltage across R_2 to rise as well."

Since the surface structure of causal arguments is almost always directed, it can be hard to distinguish between directed versus non-directed connectives. We take as evidence for non-directed connectives statements of the form "A causes B" and "B causes A" by a subject about the same situation where no signficant state change has occurred.

Conjecture 1: Of these four, the class most prototypical in human causal reasoning is explicit mechanisms with directed connectives.

For instance, subject OB was asked "If air temperature goes down, what happens to the air pressure (assuming a closed room)?" He replied:

> "As the air temperature goes down, the particles move less quickly, so it lowers air pressure."

Here he reasons that the drop in air temperature means a decrease in the speed of the air molecules, which causes a drop in pressure[5].

Conjecture 2: Two exceptions to Conjecture 1 may be experts and young children. While we suspect explicit mechanims and directed primitives are typical for human causal reasoning, we think there are two clear exceptions:

Experts: Expert models in certain domains, such as electronics, appear to be non-directed. Further, experts know how to use constraint equations and conservation laws, and therefore can reason non-mechanistically. We believe experts still use directed connectives when appropriate (such as arguments about force transmission), but also have other options.

Young Children: According to Piaget (1960), very young children lack notions of directed mechanisms; not until about 8 years old do children show fully mechanistic reasoning. Piaget's

---

[5] See Collins & Gentner (1983, 1986) for a more detailed treatment of mental models of evaporation.

interviews with 4–5 year–old children led him to conclude that they have synthetic, holistic understandings of causality. For example, when asked why a river flows, an adult or older child would answer in terms of the slope, or difference in height between the source and the destination. But five– and six–year–old children give very different answers: "Because people make oars. They push." or "Because there are big fish that swim." or "To make the fountains flow."

Piaget's interviews clearly show a difference between the way young children and older subjects talk about causality, However, recent research has cast doubt on his strong claim that young children lack mechanistic causality. When children are given tasks in familiar, concrete domains, and are asked to make predictions based on causal relations rather than verbally explaining them, even preschoolers show evidence of directed causal mechanisms (Baillargeon & Gelman, 1980; Bullock & Gelman, 1979; Bullock, Gelman & Baillargeon, 1982). Thus, young children are capable of directed mechanistic causality in familar domains. However, in unfamiliar domains such as evaporation or heat–flow, young children may reason non–mechanistically simply because they don't have enough domain knowledge to postulate mechanisms.

Thus, for both explicit mechanism and directed causality, there may be a U–shaped curve. We may find that both extreme novices and advanced experts show non–mechanistic reasoning about quantities, for opposite reasons. But aside from these two extreme groups, we believe commonsense causal reasoning is built around mechanisms. Indeed, we conjecture that even when people don't know the mechanism behind a change they postulate one, as in the Causal Corpus (Forbus and Gentner, 1986).

Now we turn to the choice of measurement scenarios.

Conjecture 3: The incremental scenario is the most basic of the measurement scenarios. In the most natural form of an incremental scenario, events occur in a causal chain, each event causing the next. Along with explicit mechanisms and directed connectives, this kind of sequential causality has considerable introspective appeal as a causal argument. The popularity of Rube Goldberg's elaborately tortuous causal chains is one indication of this idea's appeal. It is a robust way of reasoning about mechanistic causality. There is evidence that it is learned very early, at least for familiar devices (Bullock, Gelman & Baillargeon, 1982).

Examples of the use of incremental scenario occur in people's reasoning about flow systems such as electricity (Gentner & Gentner, 1983). Here subject DDD is answering the question: "Why do electrical plugs have two prongs?"

> "...little negative electrons get forced into that one prong — like suddenly there is this new space for them to go into and they have been lying in wait in my wall, waiting for this prong to come in, and they go into that one prong, and through my light bulb...And then it makes the light...it zips right down the other side of the big loop and you have a current going, and it makes that little light which is why we tricked it into my plugs. That must be why there are two plugs [prongs]. It's a differential and you need two things for there to be a difference."

Notice that the subject followed the electrons on their path from the wall to the light and back. A simpler answer is that current flows because of a voltage difference across the two prongs. But although the subject alludes to these quantities at the end of the passage, her natural approach to the question is to reason incrementally.

As another example, Subject OB is told about a pot of water sitting in a closed room and is asked "As the water temperature goes up, what do you think would be the effect on evaporation rate?" He states:

> "The water temperature going up – that means that the particles are moving faster, so they're more likely to escape, and so therefore we get a higher evaporation rate."

Conjecture 4 (Corollary to Conjecture 3): Incremental causality is so psychologically natural that people often rely on it even inapplicable. In some situations the incremental scenario will lead to incorrect conclusions. In this protocol, for example, subject CL is asked to describe what happens in the simple system shown in Figure 1.

> "All right, as the water emerges from the pump it flows at constant velocity through pipes of equivalent volume... As the pipe constricts, the flow of the water becomes slower... As it emerges from that constriction... there's a surge as the pipe expands; there's a surge in the velocity of the water, and velocity becomes slightly greater than the initial velocity

This response illustrates the difficulty novice reasoners typically have with steady-state systems. Like most novices, CL does not understand Bernoulli's principle: his expectation is that the water will slow down in the narrow pipes, whereas in fact the opposite will occur. In his reasoning, the water starts at the pump and heads into the system, encountering obstacles along the way. It is like turning a piece of stuff loose at the start of a toboggan run and watching its progress. This incremental argument leads to problems because it leads one to believe that the pieces of stuff "pile up" against each other when they reach a constriction, and thus slow down.

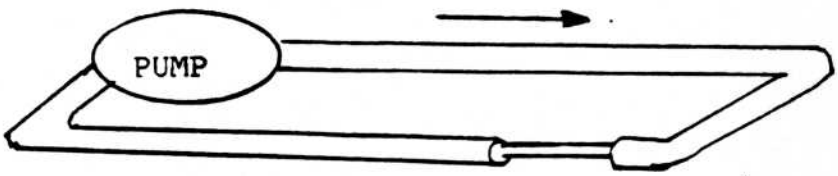

Figure 1 – A simple fluid system
The fluid system shown below consists of a pump and a constriction. Subjects are asked to reason about what happens to various quantities at different parts of the circuit.

Abandoning the incremental model for a steady state model reduces the chance of this plausible error.

The "naive incrementalism" illustrated in this protocol seems to apply to other domains as well. In electricity, for example, novices typically maintain that any of several quantities – voltage, current, power, force, energy, or velocity of electrons – is large at the start of the circuit, just after the battery, and small at the end of the circuit, just before the battery (Gentner & Gentner, 1983).

Conjecture 5: Differential Scenarios are also common in reasoning about quantities. As an example, Subject OB, asked whether an increase in water temperature will affect air pressure, says:

> "OK. If the water temperature goes up, we're going to increase the evaporation rate. If we increase the evaporation rate...we're increasing the amount of water in the air and therefore the air pressure will go up at an increasing rate. Again, it was increasing anyway, so now it's increasing a little faster."

4. Conclusions

We suggest that there are several distinct notions of causality that have psychological force. This multiplicity does not, however, render causal reasoning an inappropriate subject of study. Instead, it means that the form of our analysis must change. Studies of causal reasoning must focus on particular classes of arguments, not causal reasoning in general. While the principles obtained in this way will probably not be as general as those which are the goal of a more general analysis, we can hope that in fact we will have better success. This paper presents an example of such an analysis, examining causal arguments involving changes in physical quantities.

We have abstracted from AI research on qualitative physics three aspects of causal reasoning about quantities, *explicit versus implicit mechanisms, directed versus nondirected connectives*, and *type of measurement scenario*. We have shown where the current systems of qualitative physics lie, and shown some new directions such research can take. We have used thi analysis to draw some implications for psychology, presenting five conjectures for empirical test.

5. Acknowledgements

Several of these ideas were originally developed in collaboration with Allan Collins and Lance Rips. We also thank Renee Baillargeon and Jerry DeJong for their helpful comments. This research is supported by Office of Naval Research, Contract No. N00014-85-K-0559.

6. Bibliography

Baillargeon, R., & Gelman, R. (1980), Young children's understanding of simple causal sequences: Predictions and explanations. Presented at the meetings of the American Psychology Association, Montreal.

Bullock, M., Gelman, R. (1979), Children's assumptions about cause and effect: Temporal ordering. *Child Development, 50*, 89-96.

Bullock, M., Gelman, R. & Baillargeon, R. (1982), The development of causal reasoning. **The developmental psychology of time.** New York: Academic Press.

Bobrow, D. (Ed.), (1985), *Qualitative Reasoning about Physical Systems.* Cambridge, Mass.: MIT Press.

Bunge, M. (1979), *Causality and modern science.* New York: Dover Publications, Inc.

Collins, A.M., Warnock, E., Aiello, N., and Miller, M. (1975), Reasoning from incomplete knowledge. In D. Bobrow and A. Collins (Eds.), *Representation and understanding,* New York: Academic Press.

Collins, A.M., and Gentner, D., (1986), How people construct mental models. In D. Holland and N. Quinn (Eds.), *Cultural models in language and thought.* Cambridge, England: Cambridge University.

de Kleer, J. and Brown, J. (1984), A qualitative physics based on confluences. *Artificial Intelligence,* **24**.

DiSessa, A. (1983), Phenomenology and the evolution of intuition. In D. Gentner and A. Stevens (Eds.), *Mental Models,* Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Forbus, K. (1984), Qualitative process theory. *Artificial Intelligence,* **24**, pp 85–168.

Forbus, K. and Gentner, D. (1986), Learning physical domains: Towards a theoretical framework. In R. Michalski, J. Carbonell, and T. Mitchell (Eds.), *Machine Learning, Volume II,* Los Altos, California: Morgan Kaufmann Publishers, Inc.

Gentner, D. and Gentner, D. R. (1983), Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner and A. Stevens (Eds.), *Mental models,* Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Hayes, P. (1979), The naive physics manifesto. In D. Michie (Ed.), *Expert systems in the micro-electronic age,* Edinburgh: Edinburgh University Press.

Piaget, J. (1960), *The child's conception of physical causality.* Paterson, New Jersey: Littlefield, Adams & Co.

Simmons, R. (1983, December), Representing and reasoning about change in geologic interpretation. MIT Artificial Intelligence Lab TR–749.

Williams, B. (1984), Qualitative analysis of MOS circuits. *Artificial Intelligence,* **24**.

Weld, D. (1984), Switching between discrete and continuous process models to predict genetic activity. MIT Artificial Intelligence Lab TR–793, October.