Ann. Rev. Comput. Sci. 1988. 3: 197–232

COMMONSENSE PHYSICS: A Review

Kenneth D. Forbus

Qualitative Reasoning Group, Department of Computer Science, University of Illinois, Urbana, Illinois 61801

1. INTRODUCTION

Understanding commonsense reasoning is a central problem of Artificial Intelligence. Without a broad codification of human knowledge, and techniques for reasoning with such knowledge, our programs are doomed to remain confined to specialized areas. While expert systems have sometimes been strikingly successful in narrow, carefully defined domains, they remain brittle and hard to maintain. Natural language "front ends" are successful only when the domain of discourse is strictly limited. Robots cannot predict indirect consequences of their actions (e.g. that leaving a tool outside may cause it to rust). Really smart programs, especially those that must interact frequently with human beings, must share our common knowledge and assumptions.

What is commonsense reasoning? Sometimes it is defined only indirectly, by contrast with "expert reasoning." Some identify commonsense reasoning with default or nonmonotonic reasoning. Neither definition seems appropriate. Psychological studies of expert reasoning indicate that it relies heavily on our *mental models*, our commonsense theories of the domain of expertise (Gentner & Stevens 1983; de Kleer & Brown 1984). While experts certainly know more than novices about their domain, typically this additional knowledge is highly interconnected with the knowledge that both share. This suggests that we need to say more directly what that knowledge is. And while commonsense reasoning often involves defaults and nonmonotonic reasoning, it is hard to find areas of expertise that do not also involve defaults and nonmonotonicity. Hence such reasoning cannot be a defining property of commonsense.

Here we identify commonsense reasoning with reasoning about a par-

8756–7016/88/1115–0197$02.00

ticular collection of domains. Broadly speaking, these domains include the *physical world*, the *social world*, and the *mental world*. I focus here on reasoning about the physical world. Knowledge about the social world is typically considered by workers in natural language, since such theories are required to make sense of discourse. Knowledge about the mental world is typically the concern of workers studying theories of knowledge and action. While these areas are important, covering recent progress in reasoning about the physical world alone will strain my space limitations.

Much of the research described below is typically considered to be part of a new subfield of AI called *qualitative physics*. Qualitative physics is concerned with modeling expert reasoning as well as commonsense, and also with uncovering the ties between qualitative and traditional models of physics. We will ignore these activities here, in favor of drawing connections between qualitative physics and other areas with overlapping concerns (e.g. temporal reasoning).

The term "qualitative" has been used in many ways, sometimes just to mean "non-numerical" (e.g. "qualitative student models"). Since the term "symbolic" serves the same purpose, it seems safe to restrict "qualitative" to a more technical meaning: having to do with reasoning about continuous properties via discrete abstractions. This definition is still a bit vague, but that is to be expected given that the area is relatively new.

Qualitative physics seeks ways to represent continuous properties of the world by discrete systems of symbols. One can always quantize something continuous, but not all quantizations are equally useful. One way to state the idea is the *Relevance Principle* (Forbus 1984b): *The distinctions made by a quantization must be relevant to the kind of reasoning performed.*

The idea is simple, but few quantizations satisfy it. Rounding to fewer significant digits, replacing numbers by arbitrary intervals, using simple symbolic groups (e.g. TALL, VERY TALL), or employing fuzzy logic do not satisfy it. Representing numbers by signs (i.e. $+$ and $-$) is a useful quantization since different things tend to happen when signs change: Balls fly up and then down; different kinds of things can happen if the level of coffee in a cup is rising versus falling. Similarly, inequalities are useful because processes tend to start and stop when inequalities change: Heat flows occur when there is a temperature difference, boiling occurs when the liquid's temperature reaches its boiling point.

With the appropriate quantizations, it becomes possible to provide more abstract descriptions of state. These abstract descriptions make possible more concise descriptions of behavior. If our state parameters are elements of $\Re$, there are potentially an infinite number of states. Replacing state parameters by floating-point numbers makes the number of potential states finite, but they still number in the billions for many systems. In the

quantizations of qualitative physics there may be as few as a dozen, or a hundred, or in some cases thousands of states. Each state in a qualitative description usually corresponds to many states in a traditional description, each distinguished by having the same "meaningful behavior pattern" occurring in them.

Abstraction is a two-edged sword. While these abstract state descriptions succinctly capture possible behaviors, they tend not to prescribe exactly which behavior will occur next. By themselves they typically cannot, for we have thrown away just the information required to settle such questions. Thus qualitative simulations tend to be ambiguous. Often such ambiguous answers are good enough: If a household robot cannot imagine any way for the house to burn down as a consequence of its plan to cook supper, it can be satisfied that the plan is at least minimally safe. However, if a house fire is a possibility, more knowledge needs to be invoked. The ability of qualitative physics to represent this ambiguity explicitly is beneficial, since it provides a signal to indicate when more detailed knowledge is required. Exploring the trade-off between abstraction and precision is one of the principle themes of qualitative physics research today.

A central goal of qualitative physics is to achieve a degree of systematic coverage and uniformity far in excess of today's knowledge-based systems. In today's expert systems, knowledge is encoded about a particular domain for a particular purpose. Instead of continuing to build such systems, qualitative physics strives to create *wide-coverage*, *multi-purpose domain models*. By wide-coverage, we mean that there is some large but precisely characterizable set of systems that can be described by the domain model. It is assumed that every model for a specific system is built by instantiating appropriate elements of the domain vocabulary in appropriate ways. This will reduce the amount of hand-crafting required for new programs, and will hopefully lead to "off-the-shelf" knowledge bases.

By multi-purpose, we mean that a domain model (or a model for a specific situation) can be used for more than one inferential task. Characterizing these *styles of reasoning* is another goal of qualitative physics. These styles of reasoning include qualitative simulation, interpreting measurements, planning, comparative analysis, and others. Developing domain-independent characterizations of these styles will hopefully lead to generic algorithms that can be used as modules in a variety of larger systems.

The literature in qualitative physics has begun to grow rapidly in the last two years, and in order to focus on the new I must slight the old. While this article is self-contained, the reader who wishes to delve more deeply into the basics of the area should see Bobrow (1984) and Hobbs & Moore (1985), which contain representative samples.

2. HISTORIES, ENVISIONMENTS, AND TIME

Here I examine representations for change, focusing on temporal issues. I begin by examining the notion of *histories*, introduced by Hayes as an alternative to the situation calculus, and proceed to discuss the relationship between different forms of qualitative simulation. In discussing the notion of *envisioning* I attempt to clear up some recent confusion concerning the relationship between envisionments and histories. Finally, I survey attempts to provide general models of time and temporal reasoning.

2.1 *Histories*

Representing change is a central problem of commonsense physics. The earliest representation of changing worlds was the *situation calculus* (McCarthy & Hayes 1969). It carved up the world temporally, dividing it into *situations* of unspecified duration. A situation ends when an action occurs, and the result of taking the action on the situation gives rise to yet another situation. Thus if we have True(ON(A, B), S1) then in the situation S2 = Unstack(A, B, S1) we have True(Cleartop(B), S2). The limitations of this representation are many and well-known. It does not represent the duration of actions or indeed the duration of the situations themselves. It does not support reasoning about simultaneous or overlapping actions. But most importantly, it gives rise to the *frame problem*: Explicit axioms must be provided that state for every kind of fact whether or not it remains true after each type of action.

There have been, and continue to be, a variety of interesting attempts to solve the frame problem. One of the most productive has been the introduction of histories by Hayes (1979). Intuitively, a history is a piece of space-time, bristling with properties. A history consists of a collection of *episodes*, which serve as the spatiotemporal scope for the validity of facts associated with them. Episodes may last for an interval of time, like situations, but unlike situations they are always bounded spatially. It is assumed that for histories to interact they must intersect. This assumption provides a partial solution to the frame problem, since histories which do not touch and do not themselves have internal changes may be propagated undisturbed. The solution is not complete, since one must still determine whether histories that intersect interact, and if so, exactly how. Furthermore, some method for generating and extending histories must be specified.

Space precludes a detailed presentation of the nomenclature associated with histories, which anyway varies somewhat from researcher to researcher. The basics may be found in Hayes (1985), with some useful extensions in Forbus (1984b).

Several methods of generating histories have been devised. Hayes (1985) describes a fixed set of relationships between different types of episodes for liquids, and proposes this as a generally applicable model. Qualitative simulation can be used to produce histories directly, as Kuiper's QSIM (1986) system does (see Section 2.3), or as a side effect of quantitative simulation (Forbus 1981a). In several industrial applications, histories are specified as part of the description of a problem, as in integrated circuit fabrication (Mohammed & Simmons 1986) or hypothesized on the basis of other knowledge (Simmons 1983). One of the most complicated problems occurring in sophisticated reasoning with histories is keeping straight the justifications involved when reasoning about systems with feedback. Williams's (1986) temporal constraint propagator (TCP) is the first system that does this correctly.

2.2 *Envisionments*

A history describes a specific behavior of an object. While a history is (at least potentially) infinite, it typically consists of only a finite number of distinguishable episodes. That is, each episode can be described as an occurrence of one of a finite set of abstract *qualitative states*. This assumes there are a finite number of properties, and a finite number of values for each property, and hence only a finite number of combinations of these properties. Similarly, for any finite collection of objects we can define qualitative states that describe consistent collections of every possible distinguishable episode for each object.

Qualitative states can be defined without recourse to histories (and in fact, were developed prior to histories). Furthermore, there are several ways to compute transitions between qualitative states, depending on the nature of the physics involved. The graph formed by the collection of all qualitative states of a system and the transitions between them is called an *envisionment*. The notion of envisionment is due to de Kleer (1975). The process of constructing an envisionment, *envisioning*, was the first method of qualitative simulation.

A further distinction between envisioners is whether they start from a given initial state or from all possible states. The former are said to produce *attainable* envisionments, the latter *total* envisionments. Total envisionments are usually larger than attainable envisionments, but are more useful for certain tasks. A number of envisioners of each type have been built for different theories. NEWTON (de Kleer 1975) and FROB (Forbus 1980) both produced attainable envisionments for different kinds of motion problems. QUAL (de Kleer 1979) produced attainable envisionments for electronics, while ENVISION produced total envisionments for system-dynamics-like models (see Section 5.1). For qualitative process (QP) theory, GIZMO (Forbus 1984c) produced attainable envisionments, while QPE (Forbus 1986b) produces total envisionments.

In a correct envisionment, every possible history can be expressed as a path. Various properties of the graph correspond to important behavioral distinctions. For example, states with no transitions from them represent final states for the system, and cycles correspond to oscillations.

2.3 *Envisioning versus History Generation*

The relationship between envisionments and histories is more subtle than first suspected, and is still being explored. Some aspects are clear; for instance, Forbus (1987a) defines a *logic of occurrence* that specifies how a history may be related to an envisionment so that general behavioral constraints (such as assuming classes of behavior must or may not occur) can be enforced. Sometimes there have been simple terminological confusions, such as de Kleer & Brown (1984) calling their qualitative states "episodes," or Kuipers (1986) calling his account of history generation a "deeper semantics" for envisioning. Other aspects, however, are genuinely problematic and have become fertile areas of research.

Originally, de Kleer (1979, 1984a) claimed that, just as every history corresponds to a path through the envisionment, so every path through the envisionment must correspond to a physically realizable history. Kuipers (1986) shows this is incorrect. The counterexample he uses is shown in Figure 1 (this envisionment was generated with QPE). The parameter Z is a function of position, and should be compared with Z', but is otherwise unconstrained. By declaring the comparison between Z and Z' as interesting, we will cause a state transition to occur whenever the relationship between them changes. There are other transitions that will occur owing to the way motion and acceleration are modeled (see Forbus 1984b for details).

To generate a history from an envisionment, begin by selecting a start state. That state forms what occurs at the first episode in the history, the duration of the episode being the duration of the corresponding qualitative state (i.e. either an interval or instant). If there are no transitions from the chosen state, then that episode is the end of the history. If there are, select one of the transitions as representing what actually occurs. Then continue as before, starting from the state resulting from the transition.

Carrying out this procedure on the envisionment of Figure 1 reveals a variety of possible histories. For example, the sequence of states S_1, S_4, S_7, S_{10}, S_{13}, S_{16}, S_{19}, S_{22} corresponds to a legal history, as does S_3, S_6, S_9, S_{12}, S_{15}, S_{18}, S_{21}, S_{24}. Other legal histories correspond to variations

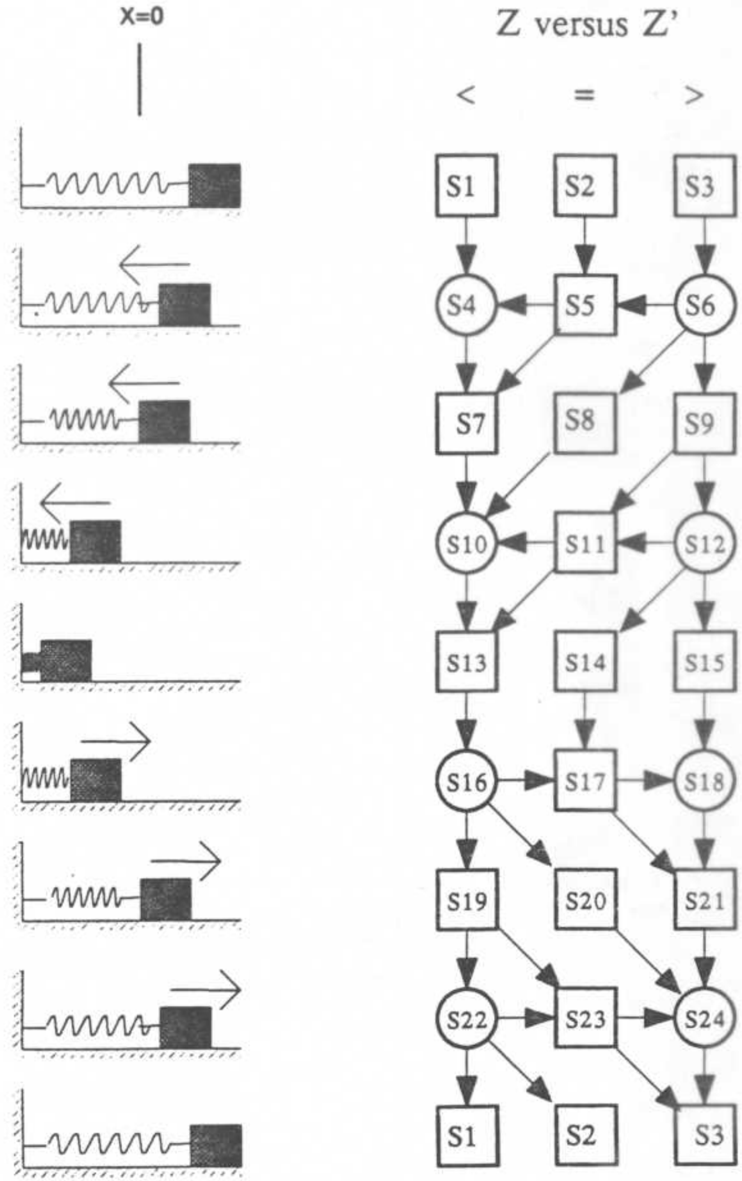

Figure 1 Generating histories from envisionments can be difficult. An envisionment for a modified spring-block oscillator is shown. The modification consists of an extra parameter Z which is a function of X and is compared with an arbitrary constant Z'. Each row is labeled with a picture indicating the general position and velocity of the block in the states of that row. Each column indicates the relationship Z has with Z' in those states. Arrows denote locally consistent transitions between states. Circles indicate states that last over an interval, while squares indicate states lasting only for an instant.

of these where Z changes in its relationship to Z' within the range of variation for X. For example, the sequence S_3, S_6, S_8, S_{10}, S_{13}, S_{16}, S_{20}, S_{24} corresponds to the case where Z equals Z' when X equals zero.

All of the histories mentioned so far are legitimate. But consider again the transitions from, say, S_6. Each time around the cycle, one of these transitions must be chosen. In the algorithm specified, which corresponds to the original de Kleer claim, each such choice is independent. Thus we are free to choose another transition the next time we reach S_6, which will give us an illegitimate history. The problem can arise even on a single cycle; the sequence S_3, S_6, S_8, S_{10}, S_{13}, S_{16}, S_{17}, S_{18}, S_{21}, S_{24} is inconsistent because the S_6, S_8, S_{10} subsequence assumes $Z = Z'$ when $X = \mathrm{ZERO}$, while the S_{16}, S_{17}, S_{18}, S_{21} is based on the assumption that Z reaches Z' before X reaches ZERO. The choices are not in fact independent, and treating them as such can lead to incorrect predictions.

In this simple case, the solution seems clear: Each choice of transition implies additional information about the functional relationship between X and Z. For example, assuming that the transition from S_6 to S_8 occurs "fixes" a point on the (implicit) graph defining their relationship: in particular, $Z = Z'$ when $X = \mathrm{ZERO}$. (Assuming that one of the other transitions occurs requires introducing a new constant related to either X or to Z, but the principle is the same). These constraints must then be respected in successive choices. For example, choosing the transition from S_{12} to S_{11} forces the later transition of S_{16} to S_{17}. However, it is not straightforward to generalize this technique to all situations.

To summarize: With no information, we can get incorrect predictions. If we had a fully specified correct quantitative model, there would be no ambiguity and hence we would always get correct histories. The open research question right now is, just how much information, and in what form, suffices to generate histories correctly from envisionments?

This problem arises even without envisionments; direct history generation must also take into account constraints imposed by earlier choices. In QSIM, for example, new named values can be introduced at every step of the computation, corresponding to the value a quantity takes on in a particular episode of the history (more on this below). Since the algorithm can introduce a new value between any two adjacent previous values, the number of possible episodes can (and does) grow exponentially without bound. This means that QSIM also produces incorrect histories. Several pruning techniques to weed out incorrect histories have been investigated, including problem-specific constraints (Lee et al 1987), algebraic manipulation (Kuipers & Chiu 1987), and quantitative knowledge (Chiu 1987), but so far these results have been mixed. [For instance, Struss (1987) points out several limitations of qualitative mathematics, such as sensitivity to the form of equations, which indicate that algebraic manipulation of qualitative equations is often unsafe.]

Both envisionment and direct history generation have their role to play in the arsenal of qualitative physics. The notion of envisionment is a superb theoretical tool, providing a simple way to think about classes of behaviors. Envisioning is a good methodological tool for qualitative model development, since it exercises domain theories in obscure cases that the model builder might otherwise ignore. But envisioning is unlikely to be the desired solution for quick on-line computation: after all, it corresponds to explicitly generating the entire problem space for some class of problems! In such cases history generation, perhaps combined with heuristics, seems to make sense. The space/time trade-offs in qualitative simulation have only begun to be explored. One can imagine compiling envisionments off-line, for example, or the envisionment of a system at a high level of abstraction being used to guide direct history generation at a lower level.

2.4 *Temporal Reasoning*

Time plays a secondary role in qualitative simulation; the relatively crude distinctions of succession and the classification of duration into instants or intervals have sufficed so far. However, there has also been much research on more general questions of temporal reasoning that is also relevant to commonsense physics. Since some of this work is reviewed by McDermott (1987), we only touch upon recent developments here. First we consider developments in temporal formalisms, and then temporal inference techniques.

Roughly, there are two main camps in AI temporal reasoning work. The first, the *numericists*, take $\Re$ as the underlying model of time, and represent intervals by segments of the real line. The primary advocates of this approach are Vere (1983), McDermott (1982), and Dean (1986). Both McDermott and Dean further use "fuzzy" bounds on duration. Their intuition is that we often do not know precise durations but can provide upper and lower bounds for how long something will take. For instance, we may know that it takes between 3 and 10 minutes to empty or fill our bathtub, between 5 and 15 minutes to actually take a bath (including drying off), and between 2 and 5 minutes to get dressed or undressed. Thus we can estimate how long it will be before the bathroom is free for someone else to use once we start to take a bath in terms of another interval, whose lower bound is the sum of the lower bounds and whose upper bound is the sum of the upper bounds. If we perform each action serially the time would be between 15 and 45 minutes.

A good planner will realize that the actions of undressing and filling the tub may be carried out in parallel, as may dressing and emptying the tub. Thus a more realistic bound would be between 11 and 35 minutes. Clearly this is a useful inference for a robot to be able to make. However, it does require knowing a lot of specific duration information to start with. Furthermore, computing the bounds in more complex examples can require a great deal of search, since there can be more than one path of intervals linking two points (Dean & McDermott 1987).

I will call the second major camp the *relationists*. The primary advocates of this approach are Allen (1984), Vilain (1986), and Ladkin (1987). The intuition behind these models is that often what we really care about is the relationships between intervals, rather than detailed information about their duration. Returning to the bathtub example, it should be clear that numerical duration information is not required to express that certain actions may be done in parallel. We need only say that both filling and undressing must be over before the bath may begin, and that both emptying and dressing must be over before someone else can take a bath. While certain inferences will require duration information—for example, figuring out whether or not five people sharing one bathroom have enough time to all take baths before a concert an hour away—relational information suffices for many inferences.

To be fair, relational systems can have computational problems as well. Often such systems are designed to make minimal temporal commitments, by describing the relationship between two intervals as a disjunction of possibilities (such as BEFORE, AFTER, MEET, DURING, etc). Processing occurs via symbolic relaxation (i.e. Waltz filtering) on triples of such relationships: For instance, if A is before B, and B is before C, then clearly A is before C. While locally such transitive inferences are cheap (they can be accomplished by table lookup), in real systems the number of them tends to grow explosively with the number of intervals in the system, with much wasted effort (Allen & Koomen 1983). However, it should be noted that this computational scheme is not inherent in the relation-based systems, and schemes involving early commitment and backup may turn out to be more tractable.

Ultimately, it will be necessary to combine something of both systems in commonsense reasoners. In constructing intricate plans, for example, one cares at first only about the rough temporal relationships between different actions and events. The relational systems provide the means of saying just what one knows, or needs to say, and no more. Precise duration estimates are irrelevant. However, when seeing if a plan will actually work, it may be necessary to resolve the ambiguities that inevitably arise when using the purely relational system. Duration estimates are appropriate for

this. However, just which aspects of each will be required and how they should interact are areas for future research.

Underlying both camps are a variety of formal models of time, and the different properties of these formalisms are the focus of much study (e.g. Hayes & Allen 1987; Ladkin 1987). To date, these efforts have not had great impact on computational schemes for temporal reasoning, mainly because the representational and inferential powers in other areas are sufficiently weak that the subtleties they address are not yet relevant to current programs. However, the future will be different. The most striking recent effort is that of Shoham (1987), who elegantly describes a formalism that captures the common properties of both the Allen and McDermott formalisms, while cleaning up various problems of each.

3. SPACE AND MOTION

Problems of spatial reasoning have tended to be neglected in qualitative physics. Even now, much of the research in this area is carried out under the aegis of vision and robotics. This is natural, since spatial reasoning is intimately connected with these areas. But there are problem-solving aspects to spatial reasoning as well that we must continue to work on if we are ever to build complete minds. Here I focus on reasoning about motion as a premiere example of such a problem. I ignore important problems such as navigation, manipulator-level planning, and layout problems, simply because they overlap significantly with robotics. Most of the projects described here are concerned with reasoning about mechanisms, such as mechanical clocks and internal combustion engines.

Motion pervades the physical world—things roll, swing, fly, gyrate, spin, and slide. To understand what happens to a tool that is placed on an incline, or to understand how mechanical clocks work, we must understand motion. As in traditional physics, we split our concerns into *qualitative dynamics* and *qualitative kinematics*. Qualitative dynamics entails organizing, representing, and reasoning with time-varying differential equations. Qualitative kinematics is the complementary study of the spatial reasoning required. Clearly, any complete account of commonsense physics must include both dynamics and kinematics. But while there has been significant progress in qualitative dynamics, there has been comparatively little progress in qualitative kinematics.

Given the importance of the problem, this lack of progress is surprising. Forbus et al (1987) attempt to explain why. They claim that "there is no purely qualitative, general-purpose kinematics" (their *poverty conjecture*). That is, the types of very weak representations used in qualitative dynamics (partial orders, monotonic functions) do not gracefully extend to higher dimensions. Instead, more powerful representations, such as symbolic algebra or numerical representations, are required as a substrate for qualitative spatial reasoning. If true, this conjecture suggests two explanations for slow progress in the area: (1) searches for a general "pure" qualitative kinematics are a waste of time, and (2) success requires multiple representations, increasing the required research effort.

They introduce the *Metric Diagram/Place Vocabulary* (MD/PV) model of spatial reasoning as an alternative to a purely qualitative representation. In this model, spatial reasoning requires two related representations. The *metric diagram* is a combination of symbolic and quantitative information that is used as an oracle for simple spatial questions. (The metric diagram is intended to serve the same role that diagrams and models play in human spatial reasoning.) The *place vocabulary* is a purely symbolic description of shape and space, grounded in the metric diagram. Since the place vocabulary is computed from the metric diagram rather than provided a priori, the qualitative representation can be chosen appropriately for the type of reasoning required.

The earliest example of this model is the FROB program (Forbus 1980, 1981a). FROB reasoned about the motion of point masses ("balls") in a 2-D world constrained by surfaces described as line segments. FROB's metric diagram consisted of symbolic descriptions of points, lines, regions, and other geometric entities containing numerical parameters. The place vocabulary was a quantization of free space, designed to maximize information about gravity and energy. These representations support a variety of inferences, including predicting the final fate of a moving object and potential collisions.

It appears that almost all spatial reasoning projects so far fit the MD/PV model. The particular form of these representations will vary with the class of problem and architecture. Bitmaps, floating-point numbers, and algebraic expressions have all been used as the quantitative component of metric diagrams. Regions of free space and configuration space have been used as the constituents of place vocabularies. The critical features of the model are that (*a*) the place vocabulary exists and (*b*) it is computed from a metric representation. These features mean that conclusions may be drawn even when little information is known (by using the place vocabulary as a substrate for qualitative spatial reasoning) and that new quantitative information, such as numerical simulations or perception, can be assimilated into the qualitative representation.

Almost all workers who tackle spatial reasoning problems invoke the equivalent of a metric diagram sooner or later. Stanfill (1983) used symbolic descriptions with numerical parameters to reason about simple pistons and bearings. Davis (1988) argues that purely qualitative representations are "too weak" to support reasoning about motion involving solid objects. Simmons's (1983) system for geological map interpretation used a combination of qualitative and quantitative simulation to produce a diagram, which was matched against geometric measurements to determine if a hypothesized sequence of geological processes could have caused the formation.

Three recent projects have focused on reasoning about mechanisms: Gelsey (1987), Joskowicz (1987), and the CLOCK project (Forbus et al 1987). I begin by examining Gelsey's program, since it is the most complete at this writing. Gelsey's program starts with a constructive solid geometry CAD description as input (his metric diagram). By computing motion envelopes and recognizing kinematic pairs, his system computes a place vocabulary consisting of regions involving interactions between parts. His program does not provide a dynamical analysis, but summarizes the results of the local kinematic analyses in terms of monotonic functions (see Section 4.2) to analyze composite systems by function composition. The program is able to analyze a significant fraction of the kinematics of an internal combustion engine, including cams and crankshaft.

The other two programs take somewhat different approaches to kinematic analysis. Faltings's (1986, 1987a,b) theory of place vocabularies for mechanisms, part of the CLOCK project, quantizes *configuration space* to serve as place vocabularies. The configuration space representation was first used in robotics for motion planning problems (Lozano-Perez 1983). The idea of configuration space is to reduce collision problems involving a moving shape and a collection of fixed shapes into a problem involving a moving point and a transformed set of obstacles. Roughly, a reference point is chosen on the moving object, and the obstacles are "grown" so that when the reference point touches the new configuration-space surface, the object will be touching the obstacle in physical space.

This representation is natural for reasoning about mechanisms for two reasons. First, connectivity is the appropriate basis for defining kinematic state, since forces change when contact relationships change. Second, Faltings observes that the important distinctions for quantizing shape must come from *pairs* of objects, rather than objects in isolation, since it is the interaction of the objects that determines if they can move or bind. In a mechanism, each part is constrained to have exactly one degree of freedom by definition. Thus a two-dimensional configuration space for each pair of parts that can touch will capture all the ways the pair of objects can interact for most mechanisms. The place vocabulary for the entire system is the combination of the place vocabularies for the components; hence the dimensionality is equal to the number of parts in the mechanism. Faltings's program is designed to work with any two-dimensional descriptions of objects in terms of line segments and arcs, and has been tested on a wide range of examples, including gears, ratchets, escapements, and a mechanical clock (see Faltings 1987b).

Joskowicz (1987) has proposed to use configuration space for analyzing mechanisms by recognition, such as determining whether or not two parts form a lower pair by inspection. (In mechanism theory a "lower pair" is formed when contact occurs along a surface; when contact involves a point or a line, the parts form a "higher pair.") He describes algorithms for doing so, and rules that could be used to recognize higher pairs. Given these descriptions, the total possible motions could be computed by constraint propagation.

The place vocabulary is used for reasoning about motion in two ways—it provides information about contact relationships, and it describes what will be reached if movement occurs in a particular direction. To use this information, qualitative frames of reference must be introduced so that forces may be imposed and motions described. Nielsen, in his part of the CLOCK project, chooses reference frames along surface normals of contacts in order to maximize dynamical information. He has also developed a qualitative *N*-dimensional vector notation, using the signs of these frames of reference. Such vectors are used for representing contact directions, forces, velocities, and other parameters. These techniques are part of his qualitative theory of rigid-body statics (Nielsen 1987), which can determine what directions an object is free to move in. This theory has been implemented and tested by determining the possible motions of gears and escapements, as well as the stability of Blocks World structures. [Shoham's (1985) formalization of freedoms, which attempts to address the same issues, is more complex and less useful than Nielsen's, since it only handles point contact and cannot represent unconstrained objects.]

4. QUANTITY AND EQUATIONS

The use of numbers to represent continuous properties and equations to represent relationships between properties is the hallmark of traditional physics. Finding less detailed, but still useful, representations for quantities and equations has been a principle focus of qualitative physics. Here I briefly review the representations currently in use, and point out recent advances.

4.1 *Quantity*

The simplest representation for numerical values are signs. For example, we might represent a voltage as being 1 if it is above some threshold, as being 0 if it equals the threshold, and -1 if it is below threshold. The

intuitive descriptions of change—"increasing," "decreasing," and "constant"—naturally correspond to the sign of a quantity's derivative. This representation is the only representation for number used by de Kleer & Brown (1984) and Williams (1984). [De Kleer's original work (1979) in electronics used sign values to stand for changes between two equilibrium states caused by a perturbation, which is a slightly different interpretation that is sometimes confused with later work.]

Other theories include signs but also allow stronger representations as well. The reason is that signs only encode a comparison with a single reference, and often more references are required (e.g. the phase of a substance being determined by the comparison between its temperature and the boiling and melting points for that substance, for the change of fluid level in a tank depending on the comparison of the relative magnitudes of flows in and out). The most common representation is inequalities, introduced in QP theory (Forbus 1981b) and used in several other systems (including Simmons 1983; Weld 1986; Kuipers 1984, 1986).

Inequalities are useful for two reasons. First, physical processes tend to start and stop when inequalities change. Flows occur when pressures or temperatures differ, for instance, and stop when they equalize. Second, they often allow sums to be disambiguated by reasoning about relative magnitudes. However, reasoning with inequalities also requires more sophisticated inference mechanisms than reasoning with signs. For example, special inference mechanisms are typically used to apply the laws of transitivity efficiently (Forbus 1984c, 1986b; Simmons 1986). It is also less intuitive to think of a numerical value as a set of statements as opposed to a discrete object. The term *quantity space* is often used to refer to the set of statements constraining the value of a number. While typically quantity spaces contain inequalities (including relationship to zero, thus including signs), recently there have been several useful extensions. For example, Simmons (1986) augments inequalities with numerical intervals and algebraic expressions, thus providing a simple way to integrate empirical bounds. I discuss other extensions below.

4.1.1 LANDMARKS AND LIMIT POINTS What should a number be compared to? Parameters representing domain-specific boundary conditions comprise one source of reference values. Examples of such *limit points* (Forbus 1981b) are the boiling temperature of a substance and the fracture stress of a material. Some comparisons are required owing to the specifics of a situation, such as a comparison between the rate of flow into and out of a container. I adopt here the terminology of Kuipers (1986) and refer to the elements of a quantity space generically as *landmark values* for the quantity, whether or not they are limit points.

Two distinct semantics have been used for landmark values in the literature. The distinction has often been misunderstood, via a type/token confusion, and I undertake to clarify it here. I call a description *temporally generic* if it refers to a class of temporal behaviors, rather than just a single behavior. A description of a single behavior I will call *temporally specific*. The script of a play is a temporally generic description, while a videotape of its performance is temporally specific. Limit points are temporally generic, as are comparisons between rates, since there are classes of situations where liquids boil and flows occur. The value of the boiling temperature at 3 p.m. is temporally specific—we are referring to a single situation, and hence a single specific value.

Most systems of qualitative physics use only temporally generic landmarks. But temporally specific landmarks can be critical for many reasoning tasks. For example, it may be crucial for a doctor to compare a patient's cholesterol level today with the specific cholesterol level last week, not just with some generic "safe" value. Kuipers QSIM generates such temporally specific landmarks. These landmarks do not correspond to "discovering" new limit points, as originally claimed. Rather, they are the equivalent of a qualitative "strip chart" that describes a specific behavior of a system. QSIM thus provides an automatic naming facility to support reasoning about temporally specific values.

Although temporally specific landmarks are essential for some inferences, they introduce a new level of computational complexity. Consider, for example, a decaying oscillation, such as a ball bouncing up and down, each time rising only some fraction of the height it reached before. Each height is a new landmark value. Thus an infinite behavior can sometimes lead to an infinite number of landmark values, as mentioned in Section 2.3.

4.1.2 ORDER OF MAGNITUDE Sometimes saying that N_1 is greater than N_2 is not enough: One may need to say that N_1 is so large compared to N_2 that N_2 may be ignored. For instance, the effect of evaporation on the level of a lake may be ignored if the dam holding it has burst. In everyday life, engineers rely on the ability to distinguish a value that is significantly out of range from a normal variation. One way to represent such information is to extend the range of comparative relationships to include *orders of magnitude*. Three such representations, FOG (Raiman 1986), O[M] (Mavrovouniotis & Stephanopolous 1987), and Davis's (1987) infinitesimal theory have been developed in qualitative physics. I begin with FOG and O[M], since they share intended use, and then describe Davis's system.

FOG introduces three new relationships, in addition to the traditional order relations. They are:

$A \ll B$ A is negligible compared to B.

$A \cong B$ A is very close to B.

$A \sim B$ A is the same order of magnitude as B.

Raiman has developed a consistent formalization that captures the intuitive meaning of these statements, using infinitesimals as a model. The effect of these relationships is to stratify values into equivalence classes, thus providing the means to say that values are very different. For example, in the DEDALE diagnosis system (Dauge et al 1987), this vocabulary is used to describe the typical relationships between values in component models.

The O[M] formalism is based on assigning labels to ranges of ratios. For example, the relationship $A \sim < B$ (read A is slightly smaller than B) is true exactly when $|A/B|$ is less than $(1+e)$, where e is a domain-specific parameter. This mapping simplifies the laws of the system and potentially allows a wide variety of quantitative information to be easily incorporated. O[M] also uses physical units to reduce inferential complexity; only parameters of the same units may be compared.

The definition of orders of magnitude relations in O[M] in terms of ranges simplifies the mapping from numerical values, a problem for which FOG provides little guidance. However, it also allows a large but finite number of negligible values to add up to something that is significant, which violates the intuitions underlying such reasoning. This cannot happen in FOG. The other relative advantages of the two systems remain to be explored.

Davis (1987) describes another formalism for order of magnitudes which, like FOG, is based on infinitesimals. He reconstructs a qualitative calculus to include infinitesimal values for both numbers and as durations of intervals. Thus he can talk about changes taking infinite (or very short) time.

4.1.3 FUZZY NUMBERS D'Ambrosio (1987) extends QP theory by using *linguistic variables*, such as "high" and "very high," defining these terms using fuzzy logic. He also annotates the qualitative equations to indicate how sensitive a result will be to a change in them. He allows such annotations to be added globally to the domain model, to specific scenarios, to specific states within a scenario, and within the scope of a single query (which could be useful in answering "what if" questions).

4.2 *Equations*

Every system of qualitative physics includes at least addition and subtraction, and multiplication is often introduced as well. However, the qualitative versions sometimes have different algebraic properties: For example signs do not form a field. Restricting ourselves to inequalities can lead to ambiguities even when an underlying field is assumed. If one only knows that A is greater than zero and B is less than zero, for instance, then the sign of $A+B$ cannot be determined. In this case knowing the relative magnitudes of A and B can provide the answer, but in general inequalities between algebraic terms are required.

In de Kleer & Brown's (1984) theory, equations constructed from the basic arithmetic operations but using only sign values are called *Confluences*. Confluences are solved by propagation of constraints, using generate-and-test when unresolvable simultaneities occur. Under certain conditions, Dormoy has shown that sets of confluences can be solved by a variant of Gaussian elimination (Dormoy & Raiman 1987). Confluences have also been used with the FOG formalism, where the comparison is made between the actual value of a parameter and its nominal value (Dauge et al 1987).

Monotonic functions are often used to approximate complicated or unknown functions with minimal commitment. From mathematics, if $y = f(x)$ then $f(x)$ is *increasing monotonic* if whenever x increases, y increases; $f(x)$ is *decreasing monotonic* if whenever x increases, y decreases. Such functions provide the weakest statement that still allows us to conclude that if we increase one quantity, the other will increase (decrease). Of course, many functions required in modeling the physical world are not monotonic. Such functions can be modeled by decomposing them into monotonic segments. Providing a framework for explicitly describing the assumptions underlying this decomposition is one of the roles played by ontology in qualitative physics (see Section 5).

Several variations of this idea are used in qualitative physics and mathematics. For instance, Kuipers uses $M^+(x,y)$ to denote an anonymous increasing monotonic connection between x and y, and $M^-(x,y)$ similarly to denote a decreasing function. QP theory allows partial specification of monotonic functions through *qualitative proportionalities*. Formally, $y \propto_{Q+} x$ indicates $y = f(\ldots, x, \ldots)$, where f is increasing monotonic in its dependence on x. Similarly, $y \propto_{Q-} x$ indicates that the function involved is decreasing monotonic in x. Qualitative proportionalities allow knowledge of equations to be distributed through a domain model, and automatically composed when needed. For example, sometimes the fluid resistance of a flow path matters and sometimes it doesn't. When it does a description adding the right $\propto_{Q-}$ can be activated to make this dependence explicit. However, using this representation requires additional inferential work, in the form of explicit closed-world assumptions.

Often additional specification of functions is allowed. *Correspondences* allow inequality information to propagate across monotonic functions. Intuitively, a correspondence fixes a point on the curve relating two (or more) parameters. An example of a correspondence is the fact that the force exerted by a spring is zero when its length equals its rest length. That fact, combined with the fact that the force increases monotonically with length, suffices to infer that if we stretch the spring beyond its rest length it will exert a positive force. Naming functions allows inequality information to be propagated across distinct individuals. Essentially, a correspondence can be established between the corresponding parameters of two distinct individuals that the function is applied to. For example, if we see two identical glasses, partially filled with water, and the level of water is the same in both, we assume the amount of water is the same as well. Details of these variations can be found in Forbus (1984c) and Kuipers (1984).

5. ONTOLOGY

Ontological choices are central to qualitative physics because, along with space and time, they provide the organizational structure for everything else. Continuous properties are properties of something, and equations hold as a result of that. Developing the appropriate ontology is usually the most difficult part of formalizing a domain.

If we are to build a complete qualitative physics, one that covers the breadth and depth of our commonsense knowledge of the physical world, we must discover and utilize common abstractions. Generating ad hoc models for each situation is impractical and unreliable. Two such ontological abstractions, *devices* and *processes*, have been widely used in qualitative physics, and I review them briefly first. Then I discuss recent progress in reasoning about substances.

5.1 *The Device Ontology*

System dynamics (Shearer et al 1971) is an engineering methodology that provides a common set of abstractions which encompass a variety of domains, including many electrical, thermal, mechanical, and acoustical systems. This modeling paradigm has been widely used in qualitative physics as well, the principle advocates being de Kleer (1979, 1984a), de Kleer & Brown (1984), and Williams (1984). These theories replace the quantitative equations of system dynamics with qualitative equations, and have developed new inference techniques for exploiting these descriptions.

The basic idea is to view a system as constructed from a collection of *devices*, such as transistors and resistors. The behavior of a device is specified by internal laws, often decomposed into distinct states or operating regions. Each device has some number of *ports*, and all interaction between devices occurs through these ports. To model a particular system, one builds a network of devices. The device network is then analyzed by using the combined equations from the devices and interconnections, either by constraint propagation or symbolic relaxation.

The device ontology has two advantages. First, the fixed network topology provides a substrate for efficient computations. All references within laws are strictly local, and hence resolving them is straightforward. Second, since system dynamics is a widely-used traditional engineering methodology, there are generally accepted standards for structural descriptions (i.e. schematics) and standard quantitative models for many domains that can be used as a starting point for creating qualitative models. The translation of such quantitative to qualitative models is not trivial, since new device states may have to be introduced (see de Kleer & Brown 1984 for details).

The device ontology has two serious limitations. First, it provides no guidance for the construction of the network model itself. In some domains (i.e. electronics) this is not a problem, but in many domains it is. A simple metal block, for instance, can be modeled as either a mass, spring, or damper, depending on the conditions around it (see Shearer et al 1971 for details). The second disadvantage is that many phenomena do not fit neatly within the device ontology at all. Two examples include unconstrained motion and phase changes. In some engineering applications such limitations can often be overcome, but it is hard to see how to do so for general commonsense physics.

5.2 *Processes*

Informally, people often describe changes in the physical world in terms of *processes*. Examples include motion, liquid flow, heat flow, boiling, bending, compressing, and expanding. This notion has been formalized in qualitative physics as an ontological commitment. Consider a cup under a faucet. If the faucet is turned on, there will be a *process* of liquid flow occurring from the faucet, through the fluid path formed by the space above the cup, to the cup itself. This liquid flow is not a property of either the cup, the faucet, the water, or the space above the cup. It is a new type of entity, with properties of its own, such as the rate of water flow.

In this ontology, processes like liquid flow provide the notion of mechanism for physical situations. All changes, ultimately, are assumed to be caused directly or indirectly by physical processes. A model of a domain includes a description of the kinds of objects there are, the kinds of relationships that hold between them, and the kinds of processes that can occur. To describe a specific situation, models for each of the parts and relationships are asserted. Importantly, the modeler does not directly specify what processes are possible. Instead, the process specifications in the domain model state the conditions under which they can occur, and the inference system uses these specifications to generate descriptions of the possible processes automatically.

This notion of process has been used by several researchers in qualitative physics, including Forbus (1981b, 1984b), Simmons (1983), Mohammed & Simmons (1986), Weld (1986), and Schmolze (1986). Some of these theories describe the effects of processes continuously over time (such as QP theory), while others describe processes by the net effect they have over an interval of time (i.e. Simmons and Weld). [The earliest attempts to formalize physical processes in AI preceded qualitative physics. Hendrix (1973) described processes as STRIPS-like operators augmented with equations for use in planning. Brown et al (1983) represented processes as finite-state automata, for instructional purposes. Neither representa tion used qualitative information, in the current technical sense of the term.]

The process ontology has several advantages. First, the notion of process is intuitively appealing for many domains. Objects can come into existence and vanish, for example, something that is not allowed in the device ontology. Second, processes provide a simple notion of causality by imposing a distinction between independent variables (those directly affected by processes) and dependent variables (those affected as a consequence of the independent variables changing). Section 6 examines this issue in detail. Third, it allows modeling assumptions to be explicitly represented, allowing programs to take on more of the modeling burden. [For example, the *class-wide assumptions* informally described by de Kleer & Brown (1984) can be formally expressed by combinations of individual and precondition specifications in QP theory.]

The process ontology also has disadvantages. First, there are some domains (like electronics) where the distinction between dependent and independent parameters changes according to the kind of analysis being performed. Process descriptions are hard to write for such domains. Second, the process ontology requires more inference, and the manipulation of quantified descriptions, to set up the model. And third, the process ontology has not been formally explored as much as the device ontology. There is no process-oriented engineering formalism equivalent to system dynamics, no off-the-shelf models to adapt. However, for commonsense physical reasoning (and, I would maintain, most engineering domains save electronics) the naturalness and modeling power of the process ontology makes it the appropriate choice.

5.3 *Substances*

Hayes (1985) identified two ontologies for reasoning about liquids. The first, the *piece of stuff* ontology, defines a liquid individual in terms of a particular collection of molecules. This ontology is well-behaved for many purposes: If, for example, you pour the water from one glass into another, it is clear that the same water has moved from one place to another. But for many purposes this representation is problematic. As Hayes points out, a river is still a river, despite the fact that the molecules of liquid that comprise it change continually. To handle such cases he defines the *contained liquid* ontology, which individuates liquids by containment. Thus a river is the same river because it is liquid in the same place. It vanishes only when the water is all gone. It is straightforward to extend this ontology to a *contained stuff* ontology that covers gases and solids as well (Forbus 1984c).

To date, most work in qualitative physics has used the contained stuff ontology. But contained stuffs have their limitations as well. Consider a refrigerator in operation. In the contained stuff view, nothing is moving in this system—the coolant in the compressor is simply the coolant in the compressor, exactly the same individual. But if we imagine what happens to a little piece of stuff, we see a drastically different picture. A little piece of stuff is in constant motion, flowing around the system. It expands, becomes vapor, gets compressed, and becomes liquid again. This sense of change captured by the piece of stuff ontology is crucial for many conclusions, both for commonsense and for engineering thermodynamics.

Collins & Forbus (1987) claim that reasoning about pieces of stuff in isolation is difficult because there is not enough information in a qualitative representation to establish gradients; consequently we cannot tell locally what a piece of stuff should do. They define a specialization of the piece-of-stuff ontology, called the *molecular collection* ontology, which solves this problem by using the contained stuff description as a framework. They define a piece of stuff, MC, which is large enough to have macroscopic properties yet small enough to never split up when traversing a fluid system. Where MC goes, and what happens to it in various places, is determined by the processes occurring to the corresponding contained stuffs. For any given qualitative state in the contained stuff description, their program can produce a corresponding envisionment for MC. This description can be used for several purposes, including figuring out that the refrigerator is a closed cycle and acts like a heat pump.

Schmolze (1986) includes a representation called *granules* in his proposal for a "physics for robots," a formalism intended to verify robot plans. Granules share some of the defining properties of MC, although he uses granules to reason by decomposition about the constituents of substances and identify processes. Unlike MC, his formalism requires numerical values; thus it will take more data to apply it.

Bunt (1985) has introduced *ensemble theory* to capture the properties of continuous concepts. His theory includes classical set theory as a special case, but also includes objects that do not have discrete parts. Thus ensemble theory can represent both continuous and discrete views of mass nouns, such as water. Raulefs (1987) extends ensemble theory, and proposes using it for reasoning about mixtures, flows, and chemical reactions. The approach is promising, but how it works in practice remains to be seen.

The research to date does not exhaust the possibilities for reasoning about fluid stuffs—the ability to individuate larger pieces that could split up in a fluid system would be useful for reasoning about contamination, for example.

6. CAUSALITY

Causality has been the subject of innumerable philosophical treatises and is currently spawning a similar number of AI papers. There is little agreement about what it means. Here I examine causality in the restricted domain of qualitative physics, hoping that by limiting the arena I can say more about it. I begin by focusing on causal reasoning about quantities, including the *causal ordering* proposal by Iwasaki & Simon (1986). Then I discuss Shoham's *chronological ignorance* approach, which proposes an interesting rigorous definition of causal theory.

6.1 *Causal Reasoning about Quantities*

Forbus & Gentner (1986a) analyze the notions of causal reasoning about quantities used in qualitative physics in order to isolate some distinctions that may be useful in understanding human reasoning. Roughly, these distinctions are: the temporal aspects relating cause and effect (the *measurement scenario*), whether or not the ontology contains an explicit class of mechanisms or not, and whether or not the primitives for describing equations include presuppositions about causality (*directed* versus *nondirected primitives*).

They assume that some notion of *mechanism* underlies causal reasoning (see Forbus & Gentner 1986b). In some theories, the notion of mechanism is tied to particular ontological classes. In QP theory, for instance, processes are the source of all changes. In other theories, such as de Kleer & Brown's confluence theory, the notion of mechanism arises from the interactions of the system's parts. They assume that flow of information in the model of the system directly mirrors "flow of causality" in the world.

Consider a liquid flow between two containers. In QP theory all changes would be caused by an instance of the liquid-flow process. In a confluence model the changes would arise from the interaction of the constitutive equations.

The *influences* used in QP theory (and others) to represent equations are directed primitives. Influences include qualitative proportionalities and *direct* influences (I+ and I−) needed to specify derivative relationships. We might represent the relationship between level and pressure in a contained liquid WC as:

$$\text{pressure(WC)} \propto_{Q+} \text{level(WC)},$$

where a change in level could cause a change in pressure, but not the reverse. In confluences (and others), the primitives do not carry a presupposition of causality. Thus we might say

$$\text{pressure(WC)} = \text{level(WC)}$$

but could not tell from this equation which way causality works. Notice that, at least in this case, there is a clear intuitive direction.

The confluence model relies on an input perturbation for causal analysis; the choice of input parameter provides significant constraint on the direction of propagation (which is interpreted as the direction of causation) in the system. This constraint is not quite sufficient, since it is necessary to annotate some parameters as independent, to prevent inappropriate causal deductions (de Kleer & Brown 1984, p. 73).

In theories with explicit mechanisms, what is an independent parameter is determined by what the mechanism directly affects. In QP theory, for instance, the *causal directedness hypothesis* (Forbus 1984b) expresses causality: "Changes in physical situations which are perceived as causal are due to our interpretation of them as corresponding either to direct changes caused by processes of propagation of those direct effects through functional dependencies." A process directly affects something by supplying its derivative. (Since it can supply a derivative of 0, the same notion suffices to impose causality on static situations.) In theories with implicit mechanisms, some other means of specifying independent parameters must be found. In confluences, for instance, the perturbed parameter and any annotated parameters are the independent ones.

Now we are in a position to understand Iwasaki & Simon's (1986) *causal ordering* proposal: They propose to use directed primitives, similar to qualitative proportionalities, but without associating a sign of effect (i.e. $\propto_Q$, but not $\propto_{Q+}$ or $\propto_{Q-}$). The exogenous variables of the system are used as the independent variables. Given these independent parameters, the technique of causal ordering will produce a graph of dependencies by

massaging the quantitative equations describing the system. To get the direction of change imposed by each connection, they propose to use the method of *comparative statics*, which uses quantitative information to produce a sensitivity analysis. The end result will be much the same as the graph of influences that holds for the corresponding situation in a QP model. The possibility of incorrect causal arguments seems to be avoided by detecting when the system of equations is underdetermined: It is exactly in such cases that an assumption must be made, and an external knowledge source (such as the user) can determine which assumption will lead to correct arguments.

Whether or not causal ordering is useful in analyzing a particular example depends on the availability of two things: a set of quantitative equations and knowledge about which variables are exogenous. For many circumstances equations are available, but for many simple circumstances (such as boiling) they aren't. Often the available equations are too complicated to use: A high-accuracy differential equation model of a coal-fired power plant, for instance, can be dozens of pages long. Basing the notion of causal independence on exogenous parameters limits causal ordering to creating models of specific systems in specific modes of behavior. The limitation to specific systems comes from the fact that what is exogenous often changes when a system becomes part of a larger system. Thus we cannot carry our analysis of, say, a heat exchanger, intact to the analysis of a larger system including it. The limitation to specific modes of behavior comes from the fact that the equations describing a system or object can change drastically (phase changes in fluids and turbulent versus non-turbulent flow are two examples).

While causal ordering satisfies several intuitions about commonsense reasoning, it also violates several others. For instance, many people make sophisticated causal inferences about quantities without knowing the formal laws of physics. Thus it does not explain how commonsense physics comes about. It also does not assign causality in feedback systems ("a chicken and egg problem"), although such descriptions are common in informal descriptions of how systems work. There is no reason why it couldn't; in classical simulation paradigms such "loops" in the equations are broken by delay elements (i.e. integration operators), and similar techniques can be used in qualitative equations (e.g. the QP theory notion of direct influence).

I believe that, while the techniques Iwasaki & Simon describe seem to have only limited usefulness as simulation tools, they could be valuable in the context of knowledge acquisition. Consider the problem of acquiring knowledge from textbooks. Two kinds of knowledge must be encoded. The formal aspects, the equations, must be transformed into qualitative laws. The informal aspects, the contents of the text, must be transformed into the organizational structure (typically ontological) that tells when these laws are appropriate and useful. Causal ordering and comparative statics may be useful techniques in translating the formal aspects. When combined with a system that can induce representations for the informal parts, we might be able to develop tools to semi-automatically acquire qualitative models by interacting with human experts.

6.2 *Chronological Ignorance*

Causality is a broad notion, and this breadth has lead some to despair of saying much about it in general (e.g. Hayes 1979). Others continue to seek formalisms that provide a broad account of causality. An interesting recent version of the latter is Shoham's logic of *chronological ignorance*. He proposes that causal theories take the form

$$\Phi \wedge \Theta \supset \Box\varphi.$$

Intuitively, φ is a prediction (the $\Box$ indicates that it necessarily holds), Φ are the preconditions that necessarily hold in some time before φ starts, and Θ are "disclaimers," which do not prevent the prediction as long as we are not aware of their negation. Shoham demonstrates that by postponing knowing that things have happened (hence the term "chronological ignorance"), theories of this form can provide efficient yet accurate inferences.

An aside: The accuracy issue arises from the phenomena noted by Hanks & McDermott (1986). Essentially, all previous systems of nonmonotonic logic or default reasoning allow solutions that are not intended by their designers (for details, see McDermott 1987). Perhaps one of the best signs that AI is still in its early stages is rampant disagreement among its practitioners over what is and is not reasonable. McDermott, at least for a while, appeared convinced the fault lies in logic (McDermott et al 1988). Others disagreed. A cynic might say that Hanks & McDermott simply discovered what anyone who has looked seriously at reasoning about the physical world has known all along: If you place no constraints on action and change, and tell your reasoning system to come up with all possibilities, some of its results will be rubbish. What holds for simple physics is likely to yield even more bizarre results, given the increased possibilities engendered by the addition of agency. This cynical view is too strong; simple and clear demonstrations are all too rare in AI.

Shoham's blend of rigor combined with a concern for algorithmic properties is refreshing. However, it is far from clear how completely this account captures our intuitions about causality. For instance, people are often quite comfortable with simultaneous causation, which his theory disallows. Shoham points out that theory provides little guidance for the structure of the preconditions and disclaimers. Nevertheless, his approach seems to be the most promising of the formalisms that attempt to capture the broader phenomena. Perhaps introducing the ontological notion of mechanism, which plays such a large role in intuitive explanations of causal phenomena, will provide the necessary abstractions for organizing the structure of causal theories.

7. REASONING

Here I discuss recent progress in styles of reasoning—i.e. extending the ways qualitative representations can be used. Diagnosis is not discussed, since an adequate treatment of recent work is beyond the scope of this article.

7.1 *Simulating Discontinuous Change*

Many physical systems exhibit discontinuous behavior. Diodes have operating regions with fundamentally different behaviors, and devices like latches and flip-flops are often best viewed as discrete. Most systems of qualitative physics handle such phenomena because they require all changes to be continuous. Nishida & Doshita (1987) have developed techniques for reasoning about such systems. Essentially, when faced with a discontinuous change, they introduce successive "mythical instants," each of which reduces a difference between the previous state and some new equilibrium state.

7.2 *Abstraction in Simulation*

Abstraction is crucial in dealing with complex systems. Often structural abstractions, such as considering a collection of pipes, valves, burners, and manifolds as a boiler, can be determined in advance. But abstractions of behavior must be detected dynamically, during the simulation process. For example, consider a simulation of a robot moving a mountain of sand from one place to another using a teaspoon. After looking at the simulation for just a few iterations, we realize what will happen in the long run, without waiting for each grain of sand to be moved. Weld (1986) calls this style of reasoning *aggregation*.

Weld decomposes aggregation into two problems: detecting repetitions that can safely be summarized, and generating an appropriate summary. He has outlined a general framework for solving these problems, using a QP-like process description to summarize the history generated by a discrete-process simulation. His program has successfully generated simulations and summarizations for simple molecular genetics problems, and illustrates how the techniques could be used to simulate digital logic circuits efficiently.

Another kind of abstraction ignores slow changes when reasoning about quick ones. Kuipers (1987) has developed an extension to QSIM that can perform such simulation for first-order systems. The idea is to determine the net effect of fast changes, and replace them with equivalent functional dependencies in reasoning about slow changes. This technique looks promising, assuming it can be extended to more complex behaviors.

7.3 *Qualitative Analytic Solutions*

In traditional physics, a set of equations can be solved analytically or by simulation to derive the behavior of a system. Similarly, qualitative equations are typically derived from an ontology in order to generate behavior via qualitative simulation (either envisioning or history generation; see above). Sacks (1985) has developed an analytic technique that generates qualitative descriptions from traditional equations. His program begins by generating closed-form expressions for each system parameter. Next, qualitative solutions are generated for each parameter by composing generic qualitative descriptions about how the functions in that expression behave over time. His initial QMR system could solve a variety of systems, including models of a dampened oscillator and heat dissipation.

One limitation of this approach is that most interesting equations do not have analytic solutions. Sacks's (1987a) solution is to decompose more complex systems into piecewise linear approximations, use QMR on each piece, and reconstruct the global solution from the local solutions. To support this technique, he has also developed a system that uses a hierarchy of techniques to manipulate sets of inequality constraints (Sacks 1987b).

Yip (1987) has a complementary approach to the same problem. *Phase portraits* comprise a geometric technique traditionally used in mathematics to describe complex dynamics. Yip has created a vocabulary of qualitative descriptions of phase space that formalizes the intuitions mathematicians bring to bear in understanding such portraits. Given a numerical simulation of a nonlinear system he uses this vocabulary to interpret the particular behavior, and make predictions about what the other parts of phase space must be like. These predictions will ultimately form the basis of additional numerical experiments.

7.4 *Interpreting Measurements*

Ideally, we would like our programs to gather their own data about the world. A program that worked in a power plant, for instance, should have the ability to "read the gauges" to find out what is happening inside the plant. Forbus (1986a, 1987c) calls this problem *measurement interpretation*.

His ATMI theory describes how to interpret measurements taken over a span of time in terms of qualitative states. This theory is very general. It requires domain-specific procedures for performing an initial signal/symbol translation, and (potential) envisionment must exist. An implementation has been demonstrated that works on multiple ontologies (i.e. both QP models and FROB models). However, at this writing it has only been tested on simulated data without gaps, and does not specify control strategies for handling noisy data.

7.5 *Planning*

Realistic planning requires knowing what the physical world will do, with and without the planner's actions. Unintended side effects can wreck a plan, such as failing to consider alternate paths of current flow when exploiting the steam in your shower for ironing. Sometimes physical processes must be enlisted to carry out a plan: Making coffee, for instance, requires boiling water, which in turn requires heating the water. Yet the representations used in planning tend to be quite different from those used in qualitative physics, suggesting that we may have to represent the same knowledge in two different forms.

Hogge (1987a,b) has developed an alternative: Use a *domain compiler* to turn a qualitative physics into a form usable by planners. In particular, his program takes as input a QP domain model, and produces rules suitable for a temporal planner. [The planner derives from Allen & Koomen's (1983) design, adding inference rules and other extensions; see Hogge (1987c) for details.] Given a description of liquid flow, for instance, the domain compiler produces an inference rule describing what it takes to cause a liquid flow to happen. When these rules are added to other inference rules and a specification of the actions an agent may take, the planner can create plans that involve processes as intermediaries. For example, his planner can figure out that it can fill a kettle by moving it under a faucet and turning the faucet on, and that it can get boiling water by moving the kettle to the stove, turning the stove on, and waiting.

While elegant, this approach has several limitations. The large descriptions produced by the domain compiler and the complex inferences required (especially transitivity) tend to choke the temporal planner. Compiling can also produce oversimplified models. For instance, the rules implicitly assume that any influence they impose on a quantity will actually succeed in changing that quantity. Thus a planner using these rules might assume that it can prevent an ocean linear from sinking by bailing with a teaspoon. Such simplifications are hard to avoid, since the only alternatives seem to be (*a*) constructing combinations of influences at domain compilation time, which leads to combinatorial explosions; or (*b*) increasing the amount of "run-time" QP inferencing that can occur in the planner, thus violating the domain-independent nature of the planner. At least some of these limitations could be overcome by teaming the planner with a simulator, so the initial plan could be compared against the possibilities and debugged if necessary.

7.6 *Comparative Analysis*

Many problems involve asking how some behavior would be different if something changed. Weld (1987) calls these *comparative analysis* problems. Two techniques have been developed for solving such problems. The first, *differential qualitative analysis*, was originally proposed by Forbus (1984b). However, that formulation was both incomplete and flawed. Weld has developed an elegant formalism for differential qualitative analysis, as well as algorithms for carrying it out, and has analyzed its limitations. For example, his system, CA, can figure out that if the mass in a spring-mass oscillator is made smaller, then the period of the oscillator will decrease. It does this by using the qualitative equations to figure out how the various parameters of the history will change as a function of mass, and then imagining a small mass perturbation. He has also developed a complementary technique, called *exaggeration*, which solves comparison problems by taking some value to the limit. To figure out what happens if the mass increases, for example, one might imagine the mass to be infinite. With infinite mass the period of oscillation becomes infinite (this analysis becomes complicated, and appears to require nonstandard analysis), and hence one can conclude the period will increase. Exaggeration has its limitations as well; Weld observes that it assumes the functions involved are monotonic.

7.7 *Learning*

Creating a complete qualitative physics is a Herculean task; it will become much easier if our machines can help. Several workers are tackling different aspects of this problem. Langley et al (1987) have studied various aspects of scientific discovery of physical laws. Kokar (1987) describes a methodology for determining limit points using dimensional analysis. Falkenhainer's (1985) ABACUS program uses qualitative proportionalities as an intermediate representation in inducing equations from numerical data. Mozetic (1987) describes how hierarchy can be exploited in automatically acquiring qualitative models, demonstrating his techniques with a model of the heart. Rajamoney & DeJong (1987) tackle the problem of debugging qualitative theories, providing a theoretical classification of bug types, including strategies for detecting and fixing them.

Psychological and historical evidence indicates that analogy is a powerful mechanism for learning new models (Clement 1986; Gentner 1987a,b). Falkenhainer (1987) has demonstrated that such learning is indeed computationally feasible, using his theory of *verification-based analogical learning*. His PHINEAS program learns new QP models by analogy. Given a new behavior, PHINEAS attempts to use its current domain model to explain the behavior (using the ATMI theory, described above). If ATMI succeeds, the current model explains the new behavior and nothing else occurs. But if ATMI fails, PHINEAS accesses a database of previously observed behaviors with associated explanations. He performs analogical matching on the behaviors first, to guide the transfer of a QP model from an understood domain to explain the new one. [His theory assumes Gentner's (1983) *structure mapping* theory of analogy, and uses a cognitive simulation of that theory (SME) (Falkenhainer et al 1986, 1987) as a module.] The new model, after some refinement, is then verified to see if it explains the new behavior by using ATMI once again. PHINEAS has been tested successfully on a number of examples and is undergoing further augmentation. For example, the envisionment can be viewed as a source of predictions to be further verified before accepting the new domain model. By interfacing Hogge's domain compiler (see above), PHINEAS has demonstrated a rudimentary form of experiment planning.

8. RESEARCH DIRECTIONS

While there has been great progress in qualitative physics in recent years, we are far from capturing the full power and range of human commonsense reasoning about the physical world. Here I indicate some promising lines of research, and raise some new questions.

Causality: Currently no one knows how to assign causality consistently in mixed (directed and nondirected) systems of qualitative equations. It will be interesting to see what ontological theories might provide the appropriate abstractions for use with chronological ignorance.

Reasoning: Davis (1988) argues persuasively that developing nondifferential, conservation-like arguments will greatly extend the reach of commonsense reasoning systems. For example, to conclude that a dropped glass will end up on the floor, we do not need to know exactly which part of it hits the ground first. Davis proposes to solve this problem by introducing abstract representations for space and paths. The MD/PV model will be needed to carry out these ideas, since it allows the abstract representations to communicate with each other and with quantitative data (the "final fate" computations in FROB worked this way, for a restricted class of problems).

I believe Davis is correct about the ubiquity of this kind of inference in commonsense reasoning. For example, consider how we think about an opened bottle of soda going flat. We know the CO_2 escapes via bubbles. We know practically nothing about each bubble individually, nor how many there are. But we do know that, via some path, the bubbles are escaping. If we can reach this conclusion, Weld's aggregation technique can take us the rest of the way. We need to discover the right abstractions for carrying out such inferences.

Quantities: No doubt other representations, lying between the poverty of signs and the richness of $\Re$, remain to be discovered. And no doubt there will be advances in qualitative representations for time-varying differential equations as well. But the real frontier is now partial differential equations, especially quantities that vary by space instead of time. Formalizing these "spatial quantities" will allow us to describe a vastly wider range of phenomena than at present.

I believe the problem decomposes into two parts. The first is the formalization of partial derivatives in general. While this part may have many technical obstacles, it seems likely that the current theories can be gracefully extended in this direction. The second problem appears to me to be much harder: the problem of choosing what the appropriate axes and frames of reference are.

Spatial reasoning: The metric diagram/place vocabulary model suggests that progress in spatial reasoning is strongly linked to progress in quantitative representations. Already CAD systems have been used as the substrate for qualitative reasoning, and it appears likely this trend will continue.

Integrated systems: The first systems in qualitative physics included several different representations and styles of reasoning. The focus then shifted to programs that explored in depth a particular representation and/or style of reasoning, to provide the experimental tools necessary for formalization. But the time has come to start building larger experimental systems again, using the techniques we have developed in isolation as modules in larger systems, both to evaluate how solid our progress is and to expose new research issues. Programs like PHINEAS that include several other systems as modules (QPE, ATMI, SME) are the harbinger of this trend in qualitative physics research. The wide availability of qualitative simulators like QSIM and QPE should accelerate this trend.

Another form of integration crucial to continued progress is the integration of qualitative dynamics with qualitative kinematics. A full understanding of an internal combustion engine, for instance, cannot be gleaned without understanding how physical processes and geometry interact.

Connections to perception: We view Ullman's (1985) theory of visual routines in part as a theory of human metric diagrams. Understanding

these routines could lead to improvements in qualitative kinematics, and the requirements of qualitative kinematics may in turn suggest what spatial descriptions people might be computing.

Ontology: The device ontology is more or less static; the action is in the object/process-centered theories. We need to understand more clearly the ways space interacts with individuation. I believe a superb challenge for qualitative physics is to model our commonsense understanding of the weather: what *is* a cold front, anyway? Understanding the weather will require spatial quantities, at least.

Acknowledgments

I would like to thank Johan de Kleer, Dedre Gentner, John Collins, Brian Falkenhainer, Paul Nielsen, and Ernie Davis for useful comments and discussions. Support for this endeavor has come from the Office of Naval Research (Contract No. N00014-85-K-0225, Contract No. N00014-85-K-0559) and the National Aeronautics and Space Administration (Contract No. NASA-NAG-9137).

Literature Cited

Allen, J. 1984. Towards a general model of action and time. *Artif. Intell.* 23(2): 000

Allen, J., Koomen, J. 1983. Planning using a temporal world model. *Proc. IJCAI-83, August*

Bobrow, D., ed. 1984. *Qualitative Reasoning about Physical Systems.* Cambridge, Mass: MIT Press

Brown, J., Burton, R., Zdybel, F. 1973. A model-driven question-answering system for mixed-initiative computer-assisted instruction. *IEEE Trans. Syst., Man, Cybern.*, SMC-3(2)

Bunt, H. C. 1985. The formal representation of quasi-continuous concepts. In Hobbs, R., Moore, R. (Eds.), *Formal Theories of the Commonsense World*, ed. R. Hobbs, R. Moore. Norwood, NJ: Ablex Publishing

Chiu, C. 1987. Qualitative physics based on exact physical principles. Presented at the 1st Qualitative Physics Workshop, Urbana, Illinois, May

Clement, J. 1986. Methods for evaluating the validity of hypothesized analogies. *Proc. 8th Ann. Conf. Cogn. Sci. Soc., Amherst*

Collins, J., Forbus, K. 1987. Reasoning about fluids via molecular collections. *Proc. AAAI-87, July*

D'Ambrosio, B. 1987. Extending the mathematics in qualitative process theory. *Proc. AAAI-87, Seattle, August*

Dauge, P., Raiman, O., Deves, P. 1987. Troubleshooting: when modeling is the trouble. *Proc. AAAI-87, July*

Davis, E. 1987. Order of magnitude reasoning in qualitative differential equations. NY Univ. Comput. Sci. Dept. Tech. Rep. No. 312, August

Davis, E. 1988. A logical framework for solid object physics. NY Univ. Comput. Sci. Dept. Tech. Rep. No. 245, October, 1986. To appear in *Int. J. AI in Engin*

de Kleer, J. 1975. Qualitative and quantitative knowledge in classical mechanics. TR-352, MIT AI Lab., Cambridge, Massachusetts

de Kleer, J. 1979. The origin and resolution of ambiguities in causal arguments. *Proc. IJCAI-79*, August

de Kleer, J. 1979. Causal and teleological reasoning in circuit recognition. MIT AI Lab. Tech. Rep. No. 529, September

de Kleer, J. 1984a. How circuits work. *Artif. Intell.* 24: 205–80

de Kleer, J. 1984b. Choices without backtracking. Proc. AAAI-84, Austin, Texas, August

de Kleer, J. 1986. An assumption-based truth maintenance system. *Artif. Intell.* 28: 127–62

de Kleer, J., Brown, J. 1984. A qualitative physics based on confluences. *Artif. Intell.* 24: 7–84

de Kleer, J., Williams, B. 1986. Reasoning about multiple faults. *Proc. AAAI-86, August*

Dean, T. 1986. Large-scale temporal databases. Tech. Rep. CS-86-15, Brown Univ. Dept. Comput. Sci.

Dean T., Brody, M. 1987. Incremental causal reasoning. *Proc. AAAI-87, July*

Dean T., McDermott, D. 1987. Temporal database management. *Artif. Intell.* 32: 000

Dormoy, J., Raiman, O. 1987. Assembling a device. Presented at the 1st Qualitative Physics Workshop, Urbana, Illinois, May

Falkenhainer, B. 1985. Proportionality graphs, units analysis, and domain constraints: improving the power and efficiency of the scientific discovery process. *Proc. IJCAI-85, August*

Falkenhainer, B. 1987. An examination of the third state in the analogy process: verification-based analogical learning. *Proc. IJCAI-87, August*

Falkenhainer, B., Forbus, K., Gentner, D. 1986. The structure-mapping engine. *Proc. AAAI-86, August*

Falkenhainer, B., Forbus, K., Gentner, D. 1987. The structure-mapping engine: algorithm and examples. Univ. Illinois at Urbana-Champaign, Dept. Comput. Sci. Tech. Rep. No. UIUCDCS-R-87-1361, July

Faltings, B. 1986. A theory and qualitative kinematics in mechanisms. Univ. Illinois at Urbana-Champaign, Dept. Comput. Sci. Tech. Rep. No. UIUCDCS-R-86-1274, May

Faltings, B. 1987a. Qualitative place vocabularies for mechanisms in configuration space. Univ. Illinois at Urbana-Champaign, Dept. Comput. Sci. Tech. Rep. No. UIUCDCS-R-87-1360, July

Faltings, B. 1987b. Qualitative kinematics in mechanisms. *Proc. IJCAI-87, August*

Forbus, K. 1980. Spatial and qualitative aspects of reasoning about motion. Proc. AAAI-80, Palo Alto, California, August

Forbus, K. 1981a. A study of qualitative and geometric knowledge in reasoning about motion. MIT AI Lab. Tech. Rep. No. 615, February

Forbus, K. 1981b. Qualitative reasoning about physical processes. *Proc. IJCAI-7, Vancouver, August*

Forbus, K. 1984a. An interactive laboratory for teaching control system concepts. BBN Tech. Rep. No. 5511, January

Forbus, K. 1984b. Qualitative process theory. *Artif. Intell.* 24: 000

Forbus, K. 1984c. Qualitative process theory. MIT AI Lab. Tech. Rep. No. 789, July

Forbus, K. 1985. The problem of existence. *Proc. Cogn. Sci. Soc.*

Forbus, K. 1986a. Interpreting measurements of physical systems. *Proc. AAAI-86, August*

Forbus, K. 1986b. The qualitative process engine. Tech. Rep. No. UIUCDCS-R-86-1288, December. To appear in *Int. J. AI in Engin.*

Forbus, K. 1987a. The logic of occurrence. *Proc. IJCAI-87, August*

Forbus, K. 1987b. Intelligent computer-aided engineering. *Proc. AAAI Workshop on AI in Process Engin., Columbia Univ., New York, March.* Submitted for publication

Forbus, K. 1987c. Interpreting observations of physical systems. *IEEE Trans. Syst., Man, Cybern.* SMC-17 (No. 3): 350–59

Forbus, K., Nielsen, P., Faltings, B. 1987. Qualitative kinematics: a framework. *Proc. IJCAI-87, August*

Forbus, K., Gentner, D. 1986a. Causal reasoning about quantities. *Proc. Eighth Annu. Conf. Cogn. Sci. Soc.*, Amherst, Massachusetts, August

Forbus, K., Gentner, D. 1986b. Learning physical domains: towards a theoretical framework. In *Machine Learning: An Artificial Intelligence Approach, Vol. 2*, ed. R. Michalski, J. Carbonell, T. Mitchell, pp. 311–48. Los Altos: Tioga Press

Forbus, K., Stevens, A. 1981. Using qualitative simulation to generate explanations. BBN Tech. Rep. No. 4490, March, 1981; *Proc. Third Annu. Meet.* Cogn. Sci. Soc, August

Gelsey, A. 1987. Automated reasoning about machine geometry and kinematics. *Third IEEE Conf. AI Appl.*, Orlando, Florida, February

Gentner, D., Stevens, A., eds. 1983. *Mental Models*. LEA Associates, Inc.

Gentner, D. 1983. Structure-mapping: a theoretical framework for analogy. *Cogn. Sci.* 7(2): 000

Gentner, D. 1987a. Mechanisms of analogical learning. In *Similarity and Analogical Reasoning*, ed. S. Vosniadou, A. Ortony. London: Cambridge Univ. Press

Gentner, D. 1987b. Historical shifts in the use of analogy in science. Univ. Illinois Dept. Comput. Sci. Tech. Rep. No. UIUCDCS-R-87-1389, November

Hanks, S., McDermott, D. 1986. Default reasoning, nonmonotonic logics, and the frame problem. *Proc. AAAI-86, August*

Hayes, P. 1979. The naive physics manifesto. In *Expert Systems in the Micro-Electronic Age*, ed. D. Michie. Edinburgh: Edinburgh Univ. Press

Hayes, P. 1985. Naive physics 1: ontology for liquids. In *Formal Theories of the Commonsense World*, ed. R. Hobbs, R. Moore. Norwood, NJ: Ablex

Hayes, P., Allen, J. 1987. Short time periods. *Proc. IJCAI-87, August*

Hendrix, G. 1973. Modeling simultaneous actions and continuous processes. *Artif. Intell.* 4

Hinton, G. 1979. Some demonstrations of the effects of structural descriptions in mental imagery. *Cogn. Sci.* 3(3)

Hobbs, R., Moore, R., eds. 1985. *Formal Theories of the Commonsense World*. Norwood, NJ: Ablex

Hogge, J. 1987a. Compiling plan operators from domains expressed in qualitative process theory. *Proc. AAAI-87, July*

Hogge, J. 1987b. The compilation of planning operators from qualitative process theory models. Tech. Rep. NO. UIUCDCS-R-87-1368, September

Hogge, J. 1987c. TPLAN: a temporal interval-based planner with novel extensions. Tech. Rep. No. UIUCDCS-R-87-1367, September

Hogge, J. 1987d. TIME and TPLAN user's manual. Tech. Rep. No. UIUCDCS-R-87-1366, September

Hollan, J., Hutchins, E., Weitzman, L. 1984. STEAMER: an interactive inspectable simulation-based training system. *AI Mag.*

Iwasaki, I., Simon, H. 1986. Causality in device behavior. *Artif. Intell.* 29: 000

Joskowicz, L. 1987. Shape and function in mechanical devices. *Proc. AAAI-87, July*

Kokar, M. 1987. Critical hypersurfaces and the quantity space. *Proc. AAAI-87, July*

Kosslyn, S. 1980. *Image and Mind.* Cambridge, Mass: Harvard Univ. Press

Kuipers, B. 1984. Common sense causality: deriving behavior from structure. *Artif. Intell.* 24: 169–204

Kuipers, B. 1986. Qualitative simulation. *Artif. Intell.* 29: 000

Kuipers, B. 1987. Abstraction by time-scale in qualitative simulation. *Proc. AAAI-87, July*

Kuipers, B., Chiu, C. 1987. Taming intractable branching in qualitative simulation. *Proc. IJCAI-87, August*

Ladkin, P. 1987. The completeness of a natural system for reasoning with time intervals. *Proc. IJCAI-87, August*

Langley, P., Simon, H., Bradshaw, G., Zytkow, J. 1987. *Scientific Discovery: Computational Explorations of the Creative Processes.* Cambridge, Mass: MIT Press

Lee, W. W., Chiu, C., Kuipers, B. J. 1987. Developments towards constraining qualitative simulation. Univ. Texas at Austin Artif. Intell. Lab. Tech. Rep. No. AI TR87-44, January

Lozano-Perez, T. 1983. Spatial planning: a configuration space approach. *IEEE Trans. Comput.* C-32, February

Mavrovouniotis, M., Stephanopolous, G. 1987. Reasoning with orders of magnitude and approximate relations. *Proc. AAAI-87, July*

McCarthy, J., Hayes, P. 1969. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence 4*. Edinburgh: Edinburgh Univ. Press

McDermott, D. 1987. Logic, problem solving, and deduction. *Ann. Rev. Comput. Sci.* 2: 187–229

McDermott, D. 1982. A temporal logic for reasoning about processes and plans. *Cogn. Sci.* 6

McDermott, D. 1988. The critique of pure reason. *Comput. Intell.* To appear

Mohammed, J., Simmons, R. 1986. Qualitative simulation of semiconductor fabrication. *Proc. AAAI-86, August*

Mozetic, I. 1987. The role of abstractions in learning qualitative models. *Proceedings of the 4th International Workshop on Machine Learning.* Los Altos: Morgan Kaufmann

Nielsen, P. 1987. The qualitative statics of rigid bodies. Univ. Illinois at Urbana-Champaign, Dept. Comput. Sci. Tech. Rep. No. UIUCDCS-R-87-1354, July

Nishida, T., Doshita, S. 1987. Reasoning about discontinuous change. *Proc. AAAI-87, July*

Raiman, O. 1986. Order of magnitude reasoning. *Proc. AAAI-86, August*

Rajamoney, S., DeJong, G. 1987. The classification, detection, and handling of imperfect theory problems. *Proc. IJCAI-87, August*

Raulefs, P. 1987. A representation framework for continuous dynamic systems. *Proc. IJCAI-87, August*

Sacks, E. 1985. Qualitative mathematical reasoning. *Proc. IJCAI-85, August*

Sacks, E. 1987a. Piecewise linear reasoning. *Proc. AAAI-87, July*

Sacks, E. 1987b. Hierarchical reasoning about inequalities. *Proc. AAAI-87, July*

Schmolze, J. 1986. Physics for robots. *Proc. AAAI-86, August*

Shearer, J., Murphy, A., Richardson, H. 1971. *Introduction to System Dynamics.* Reading, Mass: Addison-Wesley

Shoham, Y. 1985. Naive kinematics: one aspect of shape. *Proc. IJCAI-85, Los Angeles, August*

Shoham, Y. 1987. Temporal logics in AI: semantical and ontological considerations. *Artif. Intell.* 33(1): 89–104

Simmons, R. 1983. Representing and reasoning about change in geologic interpretation. MIT Artif. Intell. Lab. TR-749, December

Simmons, R. 1986. Commonsense arithmetic reasoning. *Proc. AAAI-86, August*