Qualitative Reasoning

Kenneth D. Forbus
Qualitative Reasoning Group
The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, IL, 60201, USA

DRAFT
Chapter for the CRC Handbook of Computer Science

Tuesday, June 04, 1996

1. Introduction

Qualitative reasoning is the area of AI which creates representations for continuous aspects of the world, such as space, time, and quantity, which support reasoning with very little information. Typically it has focused on scientific and engineering domains, hence its other name, qualitative physics. It is motivated by two observations. First, people draw useful and subtle conclusions about the physical world without differential equations. In our daily lives we figure out what is happening around us and how we can affect it, working with far less data, and less precise data, than would be required to use traditional, purely quantitative methods. Creating software for robots that operate in unconstrained environments and modeling human cognition requires understanding how this can be done. Second, scientists and engineers appear to use qualitative reasoning when initially understanding a problem, when setting up more formal methods to solve particular problems, and when interpreting the results of quantitative simulations, calculations, or measurements. Thus advances in qualitative physics should lead to the creation of more flexible software that can help engineers and scientists.

Qualitative physics began with de Kleer's investigation on how qualitative and quantitative knowledge interacted in solving a subset of simple textbook mechanics problems [de Kleer, 1977]. After roughly a decade of initial explorations, the potential for important industrial applications led to a surge of interest in the mid-1980s, and the area has been growing steadily, with rapid progress. Qualitative representations have made their way into commercial supervisory control software for curing composite materials, and the first product known to have been designed using qualitative physics techniques appeared on the market in 1994 [Shimomura et al., 1995]. Given the strong potential for industrial applications that is only starting to be realized, and its potential importance in understanding human cognition, work in qualitative modeling is likely to remain an important area in Artificial Intelligence.

This article first surveys the state of the art in qualitative representations and in qualitative reasoning techniques. The application of these techniques to various problems is discussed next.

2. Qualitative Representations

As with many other representation issues, there is no single, universal "right" or "best" qualitative representation. Instead there exists a spectrum of choices, each with their own advantages and disadvantages for particular tasks. What all of them have in common is that they provide notations for describing and reasoning about continuous properties of the physical world. Two key issues in qualitative representation are resolution and compositionality. We discuss each in turn.

Resolution concerns the level of information detail in a representation. Resolution is an issue because one goal of qualitative reasoning is to understand how little information suffices to draw useful conclusions. Low resolution information is available more often than precise information ("the car heading towards us is slowing down" versus "the derivative of the car's speed along the line connecting us is -12mph"), but conclusions drawn with low-resolution information are often ambiguous. The role of ambiguity is important: the prediction of alternate futures (i.e., "the car will hit us" versus "the car won't hit us") suggests that we may need to gather more information, analyze the matter more deeply, or take action, depending on what alternatives our qualitative reasoning uncovers. High resolution information is often needed to draw particular conclusions (i.e., a finite element analysis of heat flow within a notebook computer design to ensure that the CPU won't cook the battery), but qualitative reasoning with low-resolution representations reveals what the interesting questions are. Qualitative representations comprise one form of tacit knowledge that people, ranging from the person on the street to scientists and engineers, use to make sense of the world.

Compositionality concerns the ability to combine representations for different aspects of a phenomenon or system to create a representation of the phenomenon or system as a whole. Compositionality is an issue because one goal of qualitative physics is to formalize the modeling process itself. Many of today's AI systems are based on hand-crafted knowledge bases that express information about a specific artifact or system needed to carry out a particular narrow range of tasks involving it. By contrast, a substantial component of the knowledge of scientists and engineers consists of principles and laws that are broadly applicable, both with respect to the number of systems they explain and the kinds of tasks they are relevant for. Qualitative physics is developing the ideas and organizing techniques for knowledge bases with similar expressive and inferential power, called domain theories

The remainder of this section surveys the fundamental representations used in qualitative reasoning for quantity, mathematical relationships, modeling assumptions, causality, space, and time.

2.1 Representing quantity

Qualitative reasoning has explored tradeoffs in representations for continuous parameters ranging in resolution from sign algebras to the hyperreals. Most of the research effort has gone into understanding the properties of low-resolution representations, since the properties of high-resolution representations tend to already be well-understood due to work in mathematics.

The lowest resolution representation for continuous parameters is the status abstraction, which represents a quantity by whether or not it is "normal" [Abbott, 1988]. It is a useful representation for certain diagnosis and monitoring tasks because it is the weakest representation that can express the difference between something working and not working. The next step in resolution is the sign algebra, which represents continuous parameters as either -, + or 0, according to whether the sign of the underlying continous parameter is negative, positive, or zero. The sign algebra is surprisingly powerful: Since a parameter's derivatives are themselves parameters whose values can be represented as signs, some of the main results of the differential calculus (e.g., the mean value theorem) can be applied to reasoning about sign values [de Kleer & Brown, 1984]. This allows sign algebras to be used for qualitative reasoning about dynamics, including expressing properties such as oscillation and stability. The sign algebra is the weakest representation that supports such reasoning.

Representing continuous values via sets of ordinal relations (also known as the quantity space representation) is the next step up in resolution [Forbus, 1984]. For example, the temperature of a fluid might be represented in terms of its relationship between the freezing point and boiling point of the material that comprises it. Like the sign algebra, quantity spaces are expressive enough to support qualitative reasoning about dynamics. (The sign algebra can be modeled by a quantity space with only a single comparison point, zero.) Unlike the sign algebra, which draws values from a fixed finite algebraic structure, quantity spaces provide variable resolution because new points of comparison can be added to refine values. The temperature of water in a kettle on a stove, for instance, will likely be defined in terms of its relationship with the temperature of the stove as well as its freezing and boiling points. There are two kinds of comparison points used in defining quantity spaces. Limit points are derived from general properties of a domain as applicable to a specific situation. Continuing with the kettle example, the particular ordinal relationships used were chosen because they determine whether or not the physical processes of freezing, boiling, and heat flow occur in that situation. The precise numerical value of limit points can change over time, e.g., the boiling point of a fluid is a function of its pressure. Landmark values are constant points of comparison introduced during reasoning to provide additional resolution [Kuipers, 1986]. To ascertain whether an oscillating system is overdamped, underdamped, or critically damped, for instance, requires comparing successive peak values. Noting the peak value of a particular cycle as a landmark value, and comparing it to the landmarks generated for successive cycles in the behavior, provides a way of making this inference.

Intervals are a well-known variable-resolution representation for numerical values, and have been heavily used in qualitative reasoning. A quantity space can be thought of as partial information about a set of intervals. If we have complete information about the ordinal relationships between limit points and landmark values, these comparison points define a set of intervals that partition a parameter's value. This natural mapping between quantity spaces and intervals has been exploited by a variety of systems that use intervals whose endpoints are known numerical values to refine predictions produced by purely qualitative reasoning [Kuipers, 1994]. Fuzzy intervals have also been used in similar ways, e.g., in reasoning about control systems [Shen & Leitch, 1993].

Order of magnitude representations stratify values according to some notion of scale. They can be important in resolving ambiguities and in simplifying models because they enable reasoning about what phenomena and effects may safely be ignored in a given situation. For instance, heat losses from turbines is generally ignored in the early stages of power plant design, because the energy lost is very small relative to the energy being produced. Several stratification techniques have been used in the literature, including hyperreal numbers [Raiman, 1991], numerical thresholds [Mavrovouniotis & Stephanopoulos, 1988]. and logarithmic scales [Nayak, 1993]. Three issues faced by all these formalisms are (1) the conditions under which many small effects can combine to produce a significant effect, (2) the soundness of the reasoning supported by the formalism, and (3) the efficiency of using them. Understanding the properties of these formalisms and their tradeoffs is still an area of active research.

Although many qualitative representations of number use the reals as their basis, two other bases have been used with interesting results. One is the hyperreals, otherwise known as the infinitesimal calculus. Aside from order of magnitude representations, hyperreals have been used in modeling comparative analysis [Weld, 1990], dynamics [Davis, 1989], and time [Iwasaki et al., 1995]. The other basis for qualitative representations of number are finite algebras. One motivation for using finite algebras is that observations are often naturally categorized into a finite set of labels, i.e., very small, small, normal, large, very large. Research on such algebras is aimed at solving problems such as how increase the compositionality of such representations, e.g., how to propagate information across different resolution scales .

2.2 Representing mathematical relationships

Like number, a variety of qualitative representations of mathematical relationships have been developed, often by adopting and adapting systems developed in mathematics. Abstractions of the analytic functions are commonly used, to provide the lower resolution and compositionality desired. For example, confluences are differential equations over the sign algebra [de Kleer & Brown, 1984]. An equation such as V=IR can be expressed as the confluence

$$[V] = [I] + [R]$$

where [Q] denotes taking the sign of Q. Differential equations can also be expressed in this manner, for instance

[F] = ¶V

which is a qualitative version of F = MA (assuming M is always positive). Thus any system of algebraic and differential equations with respect to time can be described as a set of confluences.

Many of the algebraic operations taken for granted in manipulating analytic functions over the reals are not valid in weak algebras [Struss, 1988]. Since qualitative relationships are most often used to propagate information, this is not a serious limitation. In situations where algebraic solutions themselves are desirable, mixed representations that combine algebraic operations over the reals and move to qualitative abstractions when appropriate are a promising approach [Williams, 1991].

Another low-resolution representation of equations uses monotonic functions over particular ranges, i.e.,

M+(force, acceleration)

states that force depends only on the acceleration, and the function relating them is increasing monotonic [Kuipers, 1986]. Compositionality is achieved by using qualitative proportionalities [Forbus, 1984] to express partial information about functional dependency, e.g.,

force $_{Q+}$ acceleration

states that force depends on acceleration and is increasing monotonic in its dependence on acceleration, but may depend on other factors as well. Additional constraints on the function which determines force can be added by additional qualitative proportionalities, e.g.,
force $_{Q-}$ mass

states that force also depends on mass, and is decreasing monotonic in this dependence. Qualitative proportionalities must be combined via closed-world assumptions to ascertain all the effects on a quantity. Similar primitives can be defined for expressing relationships involving derivatives, to define a complete language of compositional qualitative mathematics for ordinary differential equations [Forbus, 1984]. As with confluences, few algebraic operations are valid for combining monotonic functions, mainly composition of functions of identical sign, i.e.,

$$M+(f,g) \wedge M+(g,h) \Rightarrow M+(f,h)).$$

In addition to resolution and compositionality, another issue arising in qualitative representations of mathematical relationships is causality. There are three common views on how mathematical relationships interact with the causal relationships people use in common sense reasoning. One view is that there is no relationship between them. The second view is that mathematical relationships should be expressed with primitives that also make causal implications. For example, qualitative proportionalities include a causal interpretation, i.e., a change in acceleration causes a change in force, but not the other way around. The third view is that acausal mathematical relationships give rise to causal relationships via the particular process of using them. For example, confluences have no built-in causal direction, but are used in causal reasoning by identifying the flow of information through them while reasoning with a presumed flow of causality in the physical system they model. One method for imposing causality on a set of acausal constraint equations is by computing a causal ordering [Iwasaki & Simon, 1986] that imposes directionality on a set of equations, starting from variables considered to be exogenous within the system.

Each view of causality has its merits. For tasks where causality is truly irrelevant, ignoring causality might be the best approach. To create software that can span the range of human commonsense reasoning, something like a combination of the second and third views appears necessary because the appropriate notion of causality varies. In reasoning about chemical phenomena, for instance, changes in concentration are always caused by changes in the amounts of the constituent parts, and never the other way around. In electronics, on the other hand, it is often convenient to consider voltage changes as being caused by changes in current in one part of a circuit and to consider current changes as being caused by changes in voltage in another part of the same circuit [Forbus & Gentner, 1986]. Understanding why different domain idealizations lead to different notions of causality is one of the interesting research issues in qualitative physics.

2.3 Ontology

Ontology concerns how to carve up the world, i.e., what kinds of things there are and what sorts of relationships can hold between them. Ontology is central to qualitative physics because one of its main goals is formalizing the art of building models of physical systems. A key choice in any act of modeling is figuring out how to construe the situation or system to be modeled in terms of the available models for classes of entities and phenomena. No single ontology will suffice for the span of reasoning about physical systems that people do. What is being developed instead is a catalog of ontologies, describing their properties and interrelationships and specifying conditions under which each is appropriate. While some ontologies are currently well understood, at this writing the catalog contains many gaps.

An example of ontologies will make this point clearer. Consider the representation of liquids. Broadly speaking, the major distinction in reasoning about fluids is whether one individuates fluid according to a particular collection of particles or by location. The former are called Eulerian, or piece of stuff, ontologies. The latter are called Lagrangian, or contained stuff ontologies. It is the contained stuff view of liquids we are using when we treat a river as a stable entity, even though the particular set of molecules that comprise it are changing constantly. It is the piece of stuff view of liquids we are using when we think about the changes in a fluid as it flows through a steady-state system, such as a working refrigerator. Ontologies multiply as we try to capture more of human reasoning. For instance, the piece of stuff ontology can be further divided into three cases, each with their own rules of inference: (1) molecular collections [Collins & Forbus, 1987], which describe the progress of an arbitrary piece of fluid that is small enough to never split apart but large enough to have extensive properties, (2) slices [Skorstad, 1992] which, like molecular collections, never subdivide but unlike them are large enough to interact directly with their surroundings, and (3) pieces of stuff large enough to be split into several pieces (e.g., an oil slick). Similarly, the contained stuff ontology can be further specialized according to whether or not individuation occurs simply by container (abstract contained stuffs [Hayes, 1985]) or by a particular set of containing surfaces (bounded stuffs [Kim, 1992]). Abstract contained stuffs provide a low-resolution ontology appropriate for reasoning about system-level properties in complex systems (e.g., the changes over time in a lubricating oil subsystem in a propulsion plant), while bounded stuffs contain the geometric information needed to reason about the interactions of fluids and shape in systems like pumps and internal combustion engines.

Cutting across the ontologies for particular physical domains are systems of organization for classes of ontologies. The most commonly used ontologies are the device ontology [de Kleer & Brown, 1984] and the process ontology [Forbus, 1984]. The device ontology is inspired by network theory and system dynamics. Like those formalisms, it construes physical systems as networks of devices whose interactions occur solely through a fixed set of ports. Unlike those formalisms, it provides the ability to write and reason automatically with device models whose governing equations can change over time.

The process ontology is inspired by studies of human mental models and observations of practice in thermodynamics and chemical engineering. It construes physical systems as consisting of entities whose changes are caused by physical processes. Process ontologies thus postulate a separate ontological category for causal mechanisms, unlike device ontologies, where causality arises solely from the interaction of the parts. Another difference between the two classes of ontologies is that in the device ontology the system of devices and connections is fixed over time, whereas in the process ontology entities and processes can come into existence and vanish over time. Each is appropriate in different contexts: For most purposes, an electronic circuit is best modeled as a network of devices, while a chemical plant is best modeled as a collection of interacting processes.

2.4 State, Time, and Behaviors

A *qualitative state* is a set of propositions that characterize a qualitatively distinct behavior of a system. A qualitative state describing a falling ball, for instance, would include information about what physical processes are occurring (e.g., motion downwards, acceleration due to gravity) and how the parameters of the ball are changing (e.g., its position is getting lower and its downward velocity is getting larger). A qualitative state can abstractly represent an infinite number of quantitative states: Although the position and velocity of the ball are different at each distinct moment during its fall, until the ball collides with the ground the qualitative state of its motion is unchanged.

Qualitative representations can be used to partition behavior into natural units. For instance, the time over which the state of the ball falling holds is naturally thought of as an interval, ending when the ball collides with the ground. The collision itself can be described as yet another qualitative state, and the fact that falling leads to a collision with the ground can be represented via a *transition* between the two states. If the ball has a non-zero horizontal velocity and there is some obstacle in its direction of travel, another possible behavior is that the ball will collide with that object instead of the ground. In general, a qualitative state can have transitions to several next states, reflecting ambiguity in the qualitative representations. Returning to our ball example, and assuming that no collisions with obstacles occur, notice that the qualitative state of the ball falling occurs again once the ball has reached its maximum height after the collision. If continous values are represented by quantity spaces and the sources of comparisons are limit points, then a finite set of qualitative states is sufficient to describe every possible behavior of a system. A collection of such qualitative states and transitions is called an *envisionment* [de Kleer, 1977]. Many interesting dynamical conclusions can be drawn from an envisionment. For instance, oscillations correspond to cycles of states. Unfortunately, the fixed resolution provided by limit points is not sufficient for other dynamical conclusions, such as ascertaining whether or not the ball's oscillation is damped. If comparisons can include landmark values such conclusions can sometimes be drawn, e.g., by comparing the maximum height on one bounce to the maximum height obtained on the next bounce. The cost of introducing landmark values is that the envisionment no longer need be finite; Every cycle in a corresponding fixed-resolution envisionment could give rise to an infinite number of qualitative states in an envisionment with landmarks.

A sequence of qualitative states occurring over a particular span of time is called a *behavior*. Behaviors can be described using purely qualitative knowledge, purely quantitative knowledge, or a mixture of both. If every continuous parameter is quantitative, the numerical aspects of behaviors coincide with the notion of trajectory in a state-space model. If qualitative representations of parameters are used, a single behavior can represent a family of trajectories through state space.

A closely related idea to behaviors are *histories* [Hayes, 1985]. Histories can be viewed as local behaviors, that is, how a single individual or property varies through time. A behavior is equivalent to a *global history*, that is, the union of all the histories for the participating individuals. The distinction is important for two reasons. First, histories are the dual of situations in the situation calculus: Histories are bounded spatially and extended temporally, while situations are bounded temporally and global spatially. Using histories avoids the frame problem, instead trading it for the more tractable problems of generating histories locally and determining how they interact when they intersect in space and time. The second reason is that history-based simulation algorithms can be more efficient than state-based simulation algorithms, since no committments need to be made concerning irrelevant information.

In a correct envisionment, every possible behavior of the physical system corresponds to some path through the envisionment. Since envisionments reflect only local constraints, the converse is not true; that is, an arbitrary path through an envisionment may not represent a physically possible behavior. All such paths must be tested against global constraints, such as energy conservation, to ensure their physical validity. Since the typical uses of an envisionment are to test whether an observed behavior is plausible or to propose possible behaviors, this limitation is not serious. A more serious limitation is that envisionments are often exponential in the size of the system being modeled. This means that in practice envisionments are rarely generated explicitly, and instead possible behaviors are searched in ways similar to those used in other areas of AI.

Many tasks require integrating qualitative states with other models of time, such as numerical models. Including precise information (e.g., algebraic expressions or floating-point numbers) about the endpoints of intervals in a history does not change their essential character.

2.5 Space and Shape

Qualitative representations of space and shape involve quantization, just as qualitative representations of continuous one dimensional parameters do. However, problem-independent purely qualitative spatial representations suffice for fewer tasks than in the one dimensional case, because of the increased ambiguity in higher dimensions [Forbus, Nielsen, & Faltings, 1991]. Consider for example deciding whether a protrusion can fit snugly inside a hole. If we have detailed information about their shapes we can derive an answer. If we consider a particular set of protrusions and a particular set of holes, we can construct a qualitative representation of these particular protrusions and holes that would allow us to derive whether or not a specific pair would fit, based on their relative sizes. But if we first compute a qualitative representation for each protrusion and hole in isolation, in general the rules of inference that can be derived for this problem will be very weak. Work in qualitative spatial representations thus tends to take two approaches. The first approach is to explore what aspects do lend themselves to qualitative representations. The second approach is to use a quantitative representation as a starting point and compute problem-specific qualitative representations to reason with. We summarize each in turn.

There are several purely qualitative representations of space and shape that have proven useful. Topological relationships between regions in two-dimensional space have been formalized, with transitivity inferences similar to those used in temporal reasoning identified for various vocabularies of relations [Cohn & Randall, 1992]. The beginnings of a rich *qualitative mechanics* have been developed. This includes qualitative representations for vectors using the sign of the vector's quadrent can be used to reason about possible directions of motion [Nielsen, 1988] and using relative inclination of angles to reason about linkages [Kim, 1992].

The use of quantitative representations to ground qualitative spatial reasoning can be viewed as a model of the ways humans use diagrams and models in spatial reasoning. One form of diagram representation are *occupancy arrays* that encode the location of an object by cells in a (two or three dimensional) grid [Funt, 1980][Glasgow, 1995]. These representations simplify the calculation of spatial relationships between objects (e.g., whether or not one object is above another), albeit at the cost of making object shape implicit. Another form of diagram representation uses symbolic structures with quantitative (e.g., numerical, algebraic, or interval) parameters [Forbus, 1980][Faltings, 1987][Joscowicz & Sacks, 1993]. These representations simplify calculations involving shape and spatial relationships, without the scaling and resolution problems that sometimes arise in array representations. However, they require a set of primitive shape elements that spans all of the possible shapes of interest, and identifying such sets for particular tasks can be difficult. For instance, many intuitively natural sets of shape primitives are not closed with respect to their complement, which can make characterizing free space difficult.

Diagram representations are used for qualitative spatial reasoning in two ways. The first is as a decision procedure for spatial questions. This mimics one of the roles diagrams play in human perception. Often these operations are combined with domain-specific reasoning procedures to produce an "analog" style of inference, where for instance the effects of perturbations on a structure are mapped into the diagram, the effect on the shapes in the diagram noted, and the results mapped back into a physical interpretation. The second way uses the diagram to construct a problem-specific qualitative vocabulary, imposing new spatial entities representing physical properties, such as the maximum height a ball can reach or regions of free space that can contain a motion . This is the *metric diagram/place vocabulary* model of qualitative spatial reasoning.

The best developed area in qualitative spatial representation is the representation of kinemantic mechanisms. The possible motions of objects are represented by qualitative regions in configuration space representing the legitimate positions of parts of mechanisms [Faltings, 1987]. While in principle a single high-dimensional configuration space could be used to represent a mechanism's possible motions (each dimension corresponding to a degree of freedom of a part of the mechanism), in practice a collection of configuration spaces, one two-dimensional space for each pair of parts that can interact, is used. These techniques suffice to analyze a wide variety of kinematic mechanisms [Joscowicz & Sacks, 1993].

Another important class of spatial representations concern qualitative representations of spatially distributed phenomena, such as flow structures [Kim, 1990] [Yip, 1995] and regions in phase space [Zhao, 1994][Bradley & Stolle, 1996]. These models use techniques from computer vision to recognize or impose qualitative structure upon a continuous field of information. This qualitative structure, combined with domain-specific models of how such structures tie to the underlying physics, enables them to interpret physical phenomena in much the same way that a scientist examining the data would.

2.6 Compositional Modeling, Domain Theories, and Modeling Assumptions

There is almost never a single "correct" model for a complex physical system. Most systems can be modeled in a variety of ways, and different tasks can require different types of models. The creation of a system model for a specific purpose is still something of an art. Qualitative physics has developed formalisms that combine logic and mathematics with qualitative representations to help automate the process of creating and refining models. The *compositional modeling* methodology [Falkenhainer & Forbus, 1991], which has become standard in qualitative physics, works like this: Models are created from *domain theories*, which describe the kinds of entities and phenomena that can occur in a physical domain. A domain theory consists of a set of *model fragments*, each describing a particular aspect of the domain. Creating a model is accomplished by instantiating an appropriate subset of model fragments, given some initial specification of the system (e.g., the propositional equivalent of a blueprint) and information about the task to be performed. Reasoning about appropriateness involves the use of *modeling assumptions*. Modeling assumptions are the control knowledge used to reason about the validity or appropriateness of using model fragments. Modeling assumptions are used to express the relevance of model fragments. Logical constraints between modeling assumptions comprise an important component of a domain theory.

An example of a modeling assumption is assuming that a turbine is isentropic. Here is a model fragment that illustrates how this assumption is used:

```
(defEquation Isentropic-Turbine
  ((turbine ?g ?in ?out)(isentropic ?g))
  (:= (spec-s ?in) (spec-s ?out)))
```

in other words, when a turbine is isentropic, the specific entropy of its inlet and outlet are equal. Other knowledge in the domain theory puts constraints on the predicate `isentropic`:

```
(for-all (?self (turbine ?self))
         (iff (= (nu-isentropic ?self) 1.0)
              (isentropic ?self)))
```

that is, a turbine is isentropic exactly when its isentropic thermal efficiency is 1. Even though no real turbine is isentropic, assuming that turbines are isentropic simplifies early analyses when creating a new design. In later design phases, when tighter performance bounds are required, this assumption is retracted and the impact of particular values for the turbine's isentropic thermal efficiency are explored. The consequences of choosing particular modeling assumptions can be quite complex; the fragments shown here are less than 1/4th of the knowledge expressing the consequences of assuming that a turbine is isentropic in a typical knowledge base.

Modeling assumptions can be classfied in a variety of ways. An ontological assumption describes which onotology should be used in an analysis. For instance, reasoning about the pressure at the bottom of a swimming pool is most simply performed using a contained-stuff representation, while describing the location of an oil spill is most easily performed using a piece-of-stuff representation. A perspective assumption describes which subset of phenomena operating in a system will be the subject. For example, in analyzing a steam plant one might focus on a fluid perspective, a thermal perspective, or both at once. A grain assumption describes how much detail is included in an analysis. Ignoring the implementation details of subsystems, for instance, is useful in the conceptual design of an artifact, but the same implementation details may be critical for troubleshooting that artifact. The relationships between these classes of assumptions can be complicated and domain-dependent; for instance, it makes no sense to include a model of a heating coil (a choice of granularity) if the analysis does not include thermal properties (a choice of perspective).

Relationships between modeling assumptions provide global structure to domain theories. Assumptions about the nature of this global structure can significantly impact the efficiency of model formulation, as discussed below. In principle, any logical constraint could be imposed between modeling assumptions. In practice, two kinds of constraints are the most common. The first are implications, such as one modeling assumption requiring or forbidding another. For example,

```
(for-all (?s (system ?s))
  (implies (consider (black-box ?s))
           (for-all (?p (part-of ?p ?s)) (not (consider ?p)))))
```

says that if one is considering a subsystem as a black box, then all of its parts should be ignored. Similarly,

```
(for-all (?l (physical-object ?l))
   (implies (consider (pressure ?l))
            (consider (fluid-properties ?l))))
```

states that if an analysis requires considering something's pressure, then its fluid properties are relevant.

The second kind of constraint between modeling assumptions are assumption classes. An assumption class expresses a choice required to create a coherent model under particular conditions. For example,

```
(defAssumptionClass (turbine ?self)
  (isentropic ?self)
  (not (isentropic ?self)))
```

states that when something is modeled as a turbine, any coherent model including it must make a choice about whether or not it is modeled as isentropic. The choice may be constrained by the data so far (e.g., different entrance and exit specific entropies), or it may be an assumption that must be made in order to complete the model. The set of choices need not be binary. For each valid assumption class there must be exactly one of the choices it presents included in the model.

3. Qualitative Reasoning Techniques

A wide variety of qualitative reasoning techniques have been developed which use the qualitative representations outlined above.

3.1 Model Formulation

Methods for automatically creating models for a specific task is one of the hallmark contributions of qualitative physics. These methods formalize knowledge and skills typically left implicit by most of traditional mathematics and engineering. To be sure, many models used in qualitative reasoning are still entirely hand-crafted, using system-specific laws and implicit task-specific simplifications. However, the state of the art in model formulation algorithms is advancing rapidly enough that this practice should soon become quite rare.

The simplest model formulation algorithm is to instantiate every possible model fragment from a domain theory, given a propositional representation of the particular scenario to be reasoned about. This algorithm is adequate when the domain theory is very focused and thus does not contain much irrelevant information. It is inadequate for broad domain theories, and fails completely for domain theories that include alternate and mutually incompatible perspectives (e.g., viewing a contained liquid as a finite object versus an infinite source of liquid). It also fails to take task constraints into account. For example, it is possible in principle to analyze the cooling of a cup of coffee using quantum mechanics. Even if it were possible in practice to do so, for most tasks simpler models suffice. Just how simple a model can be and remain adequate depends on the task. If I want to know if the cup of coffee will still be drinkable after an hour, a qualitative model suffices to infer that its final temperature will be that of its surroundings. If I want to know its temperature within 5% after 12 minutes have passed, a macroscopic quantitative model is a better choice. In other words, the goal of model formulation is to create the simplest adequate model of a system for a given task.

More sophisticated model formulation algorithms search the space of modeling assumptions, since they control which aspects of the domain theory will be instantiated. The model formulation algorithm of [Falkenhainer & Forbus, 1991] instantiated all potentially relevant model fragments and used an assumption-based truth maintenance system to find all legal combinations of modeling assumptions that sufficed to form a model that could answer a given query. The simplicity criterion used was to minimize the number of modeling assumptions. This algorithm is very simple and general but has two major drawbacks: (1) full instantiation can be very expensive, especially if only a small subset of the model fragments are eventually used, and (2) the number of consistent combinations of model fragments tends to be exponential for most problems. The rest of this section describes algorithms that overcome these problems.

Efficiency in model formulation can be gained by imposing additional structure on domain theories. Under at least one set of constraints, model formulation can be carried out in polynomial time [Nayak, 1994]. The constraints are (1) the domain theory can be divided into independent assumption classes, and (2) within each assumption class, the models can be organized by a (perhaps partial) simplicity ordering of a specific nature, forming a lattice of causal approximations Nayak's algorithm computes a simplest model, in the sense of simplest within each local assumption class, but does not necessarily produce the globally simplest model.

Conditions that ensure the creation of coherent models, that is, models which include sufficient information to produce an answer of the desired form, provide powerful constraints on model formulation. For example, in generating "what if" explanations of how a change in one parameter might affect particular other properties of the system, a model must include a complete causal chain connecting the changed parameter to the other parameters of interest.
This insight can be used to treat model formulation as a best-first search for a set of model fragments providing the simplest complete causal chain [Rickel & Porter, 1992]. A novel feature of this algorithm is that it also selects models at an appropriate time-scale. It does this by choosing the slowest time-scale phenomena that provides a complete causal model, since this provides accurate answers that minimize extraneous detail.

As with other AI problems, knowledge can reduce search. One kind of knowledge that experienced modelers accumulate concerns the range of applicability of various modeling assumptions, and strategies for how to reformulate when a given model proves inappropriate [Falkenhainer, 1993].

Model formulation is often an iterative process. For instance, an initial qualitative model is often generated to identify the relevant phenomena, followed by the creation of a narrowly focused quantitative model to answer the questions at hand. Similarly, domain-specific error criterion can determine that a particular model's results are internally inconsistent, causing the reasoner to restart the search for a good model. Formalizing the decision-making needed in iterative model formulation is an area of active research.

3.2 Causal Reasoning

Causal reasoning explains an aspect of a situation in terms of others, in such a way that the aspect being explained can be changed if so desired. For instance, a flat tire is caused by the air inside flowing out, either through the stem or through a leak. To refill the tire, we must both ensure that the stem provides a seal and that there are no leaks. Causal reasoning is thus at the heart of diagnositic reasoning, as well as explanation generation.

The techniques used for causal reasoning depend on the particular notion of causality used, but they all share a common structure. First, causality involving factors within a state are identified. Second, how the properties of a state contribute to a transition (or transitions) to another state are identified, to extend the causal account over time. Since causal reasoning often involves qualitative simulation, we turn to simulation next.

3.3 Simulation

The new representations of quantity and mathematical relationships of qualitative physics expand the space of simulation techniques considerably. We start by considering varieties of purely qualitative simulation, then describe several simulation techniques that integrate qualitative and quantitative information.

Understandinglimit analysis, the process of finding state transitions, is key to understanding qualitative simulation. Recall that a qualitative state consists of a set of propositions, some of them describing the values of continuous properties in the system. (For simplicity, in this discussion we will assume that these values are described as ordinal relations, although the same method works for sign representations and representations richer than ordinals.) Two observations are critical: (1) the phenomena which cause changes in a situation often depend on ordinal relationships between parameters of the situation, and (2) knowing just the sign of the derivatives of the parameters involved in these ordinal relationships suffices to predict how they might change over time. The effects of these changes, when calculated consistently, describe the possible transitions to other states.

An example will make this clearer. Consider again a pot of water sitting on a stove. Once the stove is turned on, heat begins to flow to the water in the stove because the stove's temperature is higher than that of the water. The causal relationship between the temperature inequality and the flow of heat means that to predict changes in the situation, we should figure out their derivatives, and any other relevant ordinal relationships that might change as a result. In this qualitative state, the derivative of the water's temperature is positive, and the derivative of the stove's temperature is constant. So one possible state change is that the water will reach thermal equilibrium with the stove and the flow of heat will stop. That is not the only possibility, of course. We know that boiling can occur if the temperature of the water begins to rise above its boiling temperature. That, too, is a possible transition that would end the state. Which of these transitions occurs depends on the relationship between the temperature of the stove and the boiling temperature of water.

This example illustrates several important features of limit analysis. First, surprisingly weak information (i.e., ordinal relations) suffice to draw important conclusions about broad patterns of physical behavior. Second, limit analysis with purely qualitative information is fundamentally ambiguous: It can identify what transitions might occur, but cannot by itself determine in all cases which transition will occur. Third, like other qualitative ambiguities, higher resolution information can be brought in to resolve the ambiguities as needed. Returning to our example, any information sufficient to determine the ordinal relationship between the stove temperature and boiling suffices to resolve this ambiguity. If we are designing an electric kettle, for instance, we would use this ambiguity as a signal that we must ensure that the heating element's temperature is well above the boiling point, and if we are designing a drink warmer, its heating element should operate well below the boiling point.

Qualitative simulation algorithms vary along four dimensions: (1) their initial states, (2) what conditions they use to filter states or transitions, (3) whether or not they generate new landmarks and (4) how much of the space of possible behaviors they explore. *Envisioning* is the process of generating an envisionment, e.g., generating all possible behaviors. Two kinds of envisioning algorithms have been used in practice: *attainable* envisioners produce all states reachable from a set of initial states, and *total* envisioners produce a complete envisionment. *Behavior generation* algorithms (e.g., QSIM [Kuipers, 1986]) start with a single initial state, generate landmark values, and use a variety of task-dependent constraints as filters and termination criterion (e.g., resource bounds, energy constraints).

Higher-resolution information can be integrated with qualitative simulation in several ways. One method for resolving ambiguities in behavior generation is to provide numerical envelopes to bound mathematical relationships. These envelopes can be dynamically refined to provide tigher situation-specific bounds. Such systems are called *semi-quantitative simulators* [Kuipers, 1994].

A different approach to integration is to use qualitative reasoning to automatically construct a numerical simulator that has integrated explanation facilities. These *self-explanatory simulators* [Forbus & Falkenhainer, 1990] use traditional numerical simulation techniques to generate behaviors, which are also tracked qualitatively. The concurrently evolving qualitative description of the behavior is used both in generating explanations and in ensuring that appropriate mathematical models are used when applicability thresholds are crossed. Self-explanatory simulators can be compiled in polynomial time for efficient execution, even on small computers [Forbus & Falkenhainer, 1995], or created in an interpreted environment [Iwasaki & Low, 1992], [Amador, Finkelstein & Weld, 1993].

3.4 Comparative Analysis

Comparative analysis answers a specific kind of “what if” questions, namely the changes that result from changing the value of a parameter in a situation. Given higher-resolution information information, traditional analytic or numerical sensitivity analysis methods can be used to answer these questions, but (a) such reasoning is commonly carried out by people who have neither the data nor the expertise to carry out such analyses and (b) purely quantitative techniques tend not to provide good explanations. Sometimes purely qualitative information suffices to carry out such reasoning, using techniques like *exaggeration* [Weld, 1990]. Consider for instance the effect of increasing the mass of a block in a spring-block oscillator. If the mass were infinite the block wouldn’t move at all, corresponding to an infinite period. Thus we can conclude that increasing the mass of the block will increase the period of the oscillator.

One paradox concerning comparative analysis with purely qualitative representations as an explanation for human common sense reasoning is that the set of unambiguous, sound inferences appears to be smaller than the set of common sense conclusions. It may be that human reasoning in this area is often unsound, or relies on experiential, higher-resolution information. Recently it has been suggested that comparative analysis may be an important form of inference with diagrams, since the concrete nature of diagrams avoids such ambiguities [Pisan, 1995].

3.5 Teleological Reasoning

Teleological reasoning connects the structure and behavior of a system to its goals. (By "its goals", we are projecting the intent of its designer or the observer, since purposes are often ascribed to components of evolved systems.) To describe how something works entails ascribing a function to each of its parts and to explain how these functions together achieve the goals. Teleological reasoning is accomplished by a combination of abduction and recognition. Abduction is necessary because most components and behaviors can play several functional roles [de Kleer, 1984]. A turbine, for instance, can be used to generate work in a power generation system and to expand a gas in a liquification system. Recognition is important because it explains patterns of function in a system in terms of known, commonly-used abstractions. A complex power-generation system with multiple stages of turbines and reheating and regeneration, for instance, can still be viewed as a Rankine cycle after the appropriate aggregation of physical processes involved in its operation [Everett, 1995].

3.6 Data Interpretation

There are two ways that the representations of qualitative physics have been used in data interpretation problems. The first is to explain a temporal sequence of measurements in terms of a sequence of qualitative states, the second is to create a qualitative model of phase space by interpreting the results of successive numerical simulation experiments. The underlying commonality in these problems is the use of qualitative descriptions of physical constraints to formulate compatibility constraints that prune the set of possible interpretations. We describe each in turn.

In measurement interpretation tasks, numerical and symbolic data are partitioned into intervals, each of which can be explained by a qualitative state or sequence of qualitative states. By using precomputed envisionments or performing limit analysis on-line, possible transitions between states used as interpretations can be found for filtering purposes. Specifically, if a state S1 is a possible interpretation for interval I1, then at least one transition from S1 must lead to a state which is an interpretation for the next interval. This compatibility constraint, applied in both directions, can provide substantial pruning. Additional constraints that can be applied include likelyhoods of particular states occurring, likelyhood of particular transitions occurring, and estimates of durations for particular states. Algorithms have been developed which can use all these constraints to maintain a "best" interpretation of a set of incoming measurements that operate in polynomial time [de Coste, 1991].

In phase space interpretation tasks, an analytic model or numerical simulation is used to gather information about the possible behaviors of a system, given a set of initial parameters. The geometric patterns these behaviors form in phase space are described using vision techniques to create a qualitative characterization of the behavior. Initially simulations are performed on a coarse grid, to create an initial description of phase space. This initial description is then used to guide additional numerical simulation experiments, using rules that express physical properties visually [Yip, 1991].

3.7 Planning

The ability of qualitative physics to provide predictions with low-resolution information and to determine what manipulations might achieve a desired effect makes it a useful component in planning systems involving the physical world. A tempting approach is to carry out qualitative reasoning entirely in a planner, by "compiling" the domain theory and physics into operators and inference rules [Hogge, 1987]. Unfortunately, such straightforward translations tend to have poor combinatorics. A different approach is to treat actions as another kind of state transition in qualitative simulation [Forbus, 1989]. This can be effective if qualitative reasoning is interleaved with execution monitoring [Drabble, 1993], or if used with a mixture of backward and forward reasoning with partial states [de Coste, 1994].

3.8 Spatial Reasoning

What distinguishes qualitative spatial reasoning from other forms of spatial reasoning is the extraction and explicit representation of qualitative descriptions of shape and space. Otherwise, the processing techniques are mainly borrowed from research in vision and robotics. Recently this flow has begun to reverse, with vision researchers adopting qualitative representations because they are more robust to compute from the data and are more appropriate for many tasks [Kuipers & Byun, 1991].

4. Applications of Qualitative Physics

Given that qualitative physics only began as a research enterprise in the late 1980s, it should not be surprising that there are to date few fielded applications. Applications in supervisory process control have been successful enough to be embedded in several commercial systems. Qualitative reasoning techniques were also used in the design of the Mita Corporation's DC-6090 photocopier [Shimomura et al., 1995], which came to market in 1994. Here we briefly summarize some representative application-oriented projects, in varying stages of fruition.

4.1 Monitoring and Diagnosis

Monitoring and diagnosis, while often treated as distinct problems, are in many applications deeply intertwined. Since these tasks also have deep theoretical commonalities, they are described together here. Monitoring a system requires summarizing its behavior at a level of description that is useful for taking action. Qualitative representations correspond to the natural descriptions applied by system operators and designers, and thus can help provide new opportunities for automation. An important benefit of using qualitative representations is that the concepts the software uses can be made very similar to those of people who interact with the software, thus improving the human-computer communication bandwidth. Diagnosis tasks impose similar requirements: It is rarely beneficial to spend the resources required to construct a very detailed quantitative model of the way a particular part has failed when the goal is to isolate a problem. Qualitative models often provide sufficient resolution for fault isolation. Qualitative models also provide the framework for organizing fault detection (i.e., noticing that a problem has occurred) and for working around a problem, tasks which often require quantitative information.

Order of magnitude representations can provide useful constraint in diagnosis tasks. DEDALE [Dauge et al., 1987] maps numerical measurements into order of magnitude ranges to test electronic circuits coming off assembly lines.

Operative diagnosis tasks are those where the system being monitored must continue being operated in spite of faults. One example of operative diagnosis is diagnosing engine trouble in civilian commercial aircraft. FaultFinder [Abbott et al., 1987], under development at NASA Langley Research Center, is intended to detect engine trouble and provide easily understood advice to pilots, whose information processing load is already substantial. Faultfinder prototypes compare engine data with a numerical simulation to detect the onset of a problem. A causal model, using low-resolution qualitative information (essentially, "working" versus "not working") is used to construct failure hypotheses, to be communicated to the pilot in a combination of natural language and graphics.

Many process control tasks involve monitoring. It has been demonstrated that qualitative representations can be used to provide more robust control than statistical process control in curing composite parts [LeClair et al., 1989]. This technique is called *Qualitative Process Automation* (QPA). In the early stage of curing a composite part, the temperature of the furnace needs to be kept relatively low because the part is outgasing. Keeping the furnace low during the entire curing process is inefficient, however, because lower temperatures means longer cure times. Therefore it is more productive to keep temperature low until outgassing stops, and then increase it to finish the cure process more quickly. Statistical process control methods use a combination of analytic models and empirical tests to figure out an optimal pattern of high/low cooking times. QPA incorporates a qualitative description of behavior into the controller, allowing it to detect the change in qualitative regime and control the furnace accordingly. The use of qualitative distinctions in supervisory control provided both faster curing times and higher yield rates than traditional techniques.

Many alarm conditions are specified as thresholds, indicating when a system is approaching a dangerous mode of operation or when a component is no longer behaving normally. Alarms are insufficient for fault detection, since they do not reflect the lack of normal behaviors. Experienced operators gain a "feel" for a system, and can sometimes spot potential problems long before it has become serious enough to trigger an alarm. Some of this expertise can be replicated by using a combination of causal models and statistical reasoning over historical data concerning the system in question [Doyle, 1995].

In some applications a small set of fault models can be preenumerated. The MIMIC approach [Kuipers, 1994] uses such a set of models to track the behavior of a system with a qualitative or semi-quantitative simulator. Any fault model whose simulation is inconsistent with the observed behavior can thus be ruled out.

Relying on a pre-existing library of fault models can limit the applicability of automatic monitoring and diagnosis algorithms. One approach to overcoming this limitation is to create algorithms that only require models of normal behavior. Most consistency-based diagnosis algorithms take this (i.e., GDE [de Kleer & Williams, 1986]) approach. The problem with this approach is that the ways a system can fail are still governed by natural laws, which imposes more constraint than logical consistency. Collins' *process-based diagnosis* [Collins, 1993] approach begins by using consistency-based algorithms to isolate the problem. It then uses the domain theory to generate explanations that could account for this problem, via abduction. A useful consequence of this approach is that it makes additional predictions, many of which could be tested to refine a hypothesis or are important for reasoning about safety in operative diagnosis (e.g., if a solvent tank's level is dropping because it is leaking, then where is the solvent going?).

4.2 Design

Engineering design activities are divided into conceptual design, the initial phase when the overall goals, constraints, and functioning of the artifact are established, and detailed design, when the results of conceptual design are used to synthesize a constructable artifact or system. Most computer-based design tools, such as CAD systems and analysis programs, facilitate detailed design. Yet many of the most costly mistakes occur during the conceptual design phase. The ability to reason with partial information makes qualiative reasoning one of the few technologies that provides substantial leverage during the conceptual design phase. Qualitative reasoning can also help automate aspects of detailed design.

The best example is the Mita Corporation's DC-6090 photocopier [Shimomura et al., 1995]. It is an example of a *self-maintenance machine* [Umeda et al., 1992], in which redundant functionality is identified at design time so that the system can dynamically reconfigure itself to temporarily overcome certain faults. An envisionment including fault models, created at design time, was used as the basis for constructing the copier's control software. In operation, the copier keeps track of which qualitative state it is in, so that it can produce the best quality copy it can.

In some fields experts formulate general design rules and methods, expressed in natural language. Qualitative representations can enable these rules and methods can be further formalized so that they can be used in software. In chemical engineering, for instance, design methods for distillation plants (including Seader & Westerberg and Nath & Motard) have been formalized using Qualitative Process theory and designs comperable to those in the chemical engineering research literature have been generated automatically [Sgouros, 1993].

Automatic analysis and synthesis of kinematic mechanisms has received considerable attention. Complex fixed-axis mechanisms, such as mechanical clocks, can be simulated qualitatively [Forbus, Nielsen, & Faltings, 1991], and a simplified dynamics added to produce convincing animations [Sacks & Joscowicz, 1993]. Initial forays into conceptual design of mechanisms have been made (e.g., [Faltings & Sun, 1993]), and qualitative kinematics simulation has been demonstrated to be competitive with conventional approaches in some linkage optimization problems [Yannou & Vasiliu, 1995].

Qualitative reasoning is also being used to reason about the effects of failures and operating procedures. Such information can be used in failure modes and effects analysis (FMEA).. For example, potential hazards in a chemical plant design can be identified by perturbing a qualitative model of the design with various faults and using qualitative simulation to ascertain the possible indirect consequences of each fault [Catino, 1993]. FMEA software for electrical system design is now being fielded for use in automotive design [Price et al, 1995].

4.3 Intelligent tutoring systems and learning environments

One of the original motivations for the development of qualitative physics was its potential applications in intelligent tutoring systems (ITSs) and learning environments (ILEs) [Brown et al., 1982][Hollan et al., 1984]. Qualitative representations provide a formal language for a student's mental models[Gentner & Stevens, 1983], and thus they facilitate communication between software and student . For example, [Frederiksen & White, 1990] describe a sequence of qualitative models which is designed to help students learn electronics. Student protocols can be analyzed in qualitative terms to diagnose misconceptions [de Koonig, Brenker, & Bredeweg, 1995].

Qualitative representations are being used in software for teaching plant operators and engineers. They provide a level of explanation for how things work that facilitates teaching control. For example, systems for teaching the operation of power generation plants, including nuclear plants, are under construction in various countries [Vadillo et al, 1994]. The ITSE project [Sime & Leitch, 1993] uses a hierarchy of models to help students understand a typical industrial process and design controllers for it. Qualitative representations can also help provide teaching software with the physical intuitions required to help find student's problems. For instance, in an ILE for engineering thermodynamics, qualitative representations are used to detect physically impossible designs [Forbus & Whalley, 1994].

Qualitative representations can be particularly helpful in teaching domains where quantitative knowledge is either nonexistant, inaccurate, or incomplete. For example, efforts underway to create ITSs for ecology in Brazil, to support of conservation efforts are using qualitative representations to explain how environmental conditions affect plant growth [Salles, 1995].

4.4 Cognitive Modeling

Since qualitative physics was inspired by observations of how people reason about the physical world, one natural application of qualitative physics would be cognitive simulation, i.e., the construction of programs whose primary concern is accurately modeling some aspect of human reasoning, as measured by comparison with psychological results. Suprisingly little has been done in this area. Most of the research has been concerned with modeling scientific discovery, e.g., how analogy can be used to create new physical theories and modeling scientific discovery in chemistry [Shrager & Langley, 1990]. Given the increasing importance of human-computer interaction, a better understanding of how people reason qualitatively could have substantial economic benefits.

5. Research Issues and Summary

Qualitative reasoning is rapidly moving from an area of pure research to a mature subfield with a mixture of basic and applied activities, including fielded applications. The substantial increases in available computing power, combined with the now urgent need to make software that is more articulate, suggests that the importance of qualitative reasoning will continue to grow.

Although there is a substantial research base to draw upon, there are many open problems and areas that require additional research. There are still many unanswered questions about purely qualitative representations (e.g., what is the minimum information that is required to guarentee that all predicted behaviors generated from an envisionment are physically possible?), but the richest vein of research concerns the integration of qualitative knowledge with other kinds of knowledge: numerical, analytic, teleological, etc. The work on modeling to date, while a solid foundation, is still very primitive: better model formulation algorithms, well-tested conventions for structuring domain theories, and robust methods for integrating the results of multiple models are needed. Substantial domain theories for a broad range of scientific and engineering knowledge need to be created. And finally, there are many domains where traditional mathematics has intruded, but where the amount and/or precision of the data available has not enabled it to be very successful. These areas are ripe for qualitative modeling. Examples include medicine [Bylander, et al, 1988], organizational theory [Kamps & Peli, 1995], economics [Lang et al, 1995] and ecology (c.f. [Guerrin, 1991])

6. Definining terms

comparative analysis: A particular form of "what if" question, i.e., how a physical system changes in response to the perturbation of one of its parameters.
compositional modeling: A methodology for organizing domain theories so that models for specific systems and tasks can be automatically formulated and reasoned about.
confluence: An equation involving sign values.
diagrammatic reasoning: Spatial reasoning, with particular emphasis on how people use diagrams.
domain theory: A collection of general knowledge about some area of human knowledge, including the kind of entities involved and the types of relationships that can hold between them, and the mechanisms that cause changes (e.g., physical processes, component laws, etc.). Domain theories range from purely qualitative to purely quantitative to mixtures of both.
envisionment: A description of all possible qualitative states and transitions between them for a system. Attainable envisionments describe all states reachable from a particular initial state, total envisionments describe all possible states.
landmark: A comparison point indicating a specific value achieved during a behavior, e.g., the successive heights reached by a partially elastic bouncing ball.

limit point: A comparison point indicating a fundamental physical boundary, such as the boiling point of a fluid. Limit points need not be constant over time, e..g, boiling points depend on pressure.

metric diagram: A quantitative representation of shape and space used for spatial reasoning.

model fragment: A piece of general domain knowledge that is combined with others to create models of specific systems for particular tasks.

modeling assumption: A proposition expressing control knowledge about modeling, such as when a model fragment is relevant.

physical process: A mechanism that can cause changes in the physical world, such as heat flow, motion, and boiling.

place vocabulary: A qualitative description of space or shape that is grounded in a quantitative representation.

qualitative proportionality: A qualitative relationship expressing partial information about a functional dependency between two parameters.

qualitative simulation: The generation of predicted behaviors for a system based on qualitative information. Qualitative simulations typically include branching behaviors, due to the low resolution of the information involved.

quantity space: A set of ordinal relationships that describes the value of a continuous parameter.

semi-quantitative simulation: A qualitative simulation that uses quantitative information, such as numerical values or analytic bounds, to constrain its results.

7. References

1. Abbott, K., Schutte, P., Palmer, M., & Ricks, W. (1987). Faultfinder: A diagnostic expert system with graceful degradation for onboard aircraft application. 14th International Symposium on Aircraft Integrated Monitoring Systems
2. Abbott, K. (1988). Robust Operative Diagnosis as Problem Solving in a Hypothesis Space. Proceedings of AAAI-88 369-374.
3. Amador, F., Finkelstein, A., & Weld, D. (1993). Real Time Self-Explanatory Simulation. Proceedings of AAAI-93 562-567.
4. Bradley, E. and Stolle, R. (1996). Automatic Construction of Accurate Models of Physical Systems. Annals of Mathematics of Artificial Intelligence (to appear).
5. Brown, J.W., Burton, R.R., & de Kleer, J. (1982). Pedagogical Natural Language and Knowledge Engineering Techniques in SOPHIE I, II and III. In D. Sleeman and J.S. Brown (Eds.), Intelligent Tutoring Systems. New York: Academic Press. 227-282.
6. Bylander, T.,Smith, J. W., & Svirbely, J. (1988). Qualitative Representation of Behavior in the Medical Domain. Computers and Biomedical Research, 21, 367-380.
7. Catino, C. (1993). Automated Modeling of Chemical Plants with Application to Hazard and Operability Studies. Unpublished doctoral dissertation, University of Pennsylvania: Chemical Engineering.
8. Cohn, A., & Randall, Z.(1992). An Interval Logic for Space Based on 'Connection'. Proceedings of the 10th European Conference on Artificial Intelligence. 394- 398.
9. Collins, J., & Forbus, K. (1987). Reasoning about fluids via molecular collections. Proceedings of AAAI-87 590-594.
10. Collins, J. (1993). Process- Based Diagnosis: An Approach to Understanding Novel Failures (Tech. Rep. No. 48). Evanston: Northwestern University, Institute for the Learning Sciences.
11. Dague, P., Raiman, O., & Deves, P. (1987). Troubleshooting: When Modeling Is the Trouble. Proceedings of AAAI-87 590-595.
12. Davis, E. (1989). Order of magnitude reasoning in qualitative differential equations. In D. Weld and J. de Kleer (Eds.), Readings in Qualitative Reasoning about Physical Systems. Los Altos, CA: Morgan Kaufmann.
13. de Coste, D. (1991). Dynamic Across-Time Measurement Interpretation. Artificial Intelligence, 51. 273-341.
14. de Coste, D. (1994). Goal- Directed Qualitative Reasoning with Partial States (Tech. Rep. No. 57). Evanston: Northwestern University, Institute for the Learning Sciences.
15. de Kleer, J. (1977). Multiple Representations of Knowledge in a Mechanics Problem Solver, Proceedings of IJCAI-77 299-304.
16. de Kleer, J. (1984). How Circuits Work. Artificial Intelligence, 24. 205-280.
17. de Kleer, J., & Brown, J. (1984). A Qualitative Physics Based on Confluences. Artificial Intelligence, 24. 7-83.
18. de Kleer, J., & Williams, B.C. (1986). Reasoning about Multiple Faults. Proceedings AAAI-86. 132-139.
19. de Koonig, K., Brenker, J., & Bredeweg, B. (1995). Cognitive Diagnosis Revisited. Proceedings of AI-ED 95 115-122.
20. Doyle, R. (1995). Determining the Loci of Anomalies Using Minimal Causal Models. Proceedings of IJCAI-95 1821-1827.

21. Drabble, B. (1993). Excalibur: A Program for Planning and Reasoning with Processes. Artificial Intelligence, 62(1). 1-40.
22. Dvorak, D., & Kuipers, B. (1989). Model-Based Monitoring of Dynamic Systems. Proceedings of AAAI-89.
23. Everett, J. (1995). A Theory of Mapping from Structure to Function Applied to Engineering Thermodynamics. Proceedings of IJCAI-95. 1837-1843.
24. Falkenhainer, B. (1990). A Unified Approach to Explanation and Theory Formation. Shrager & Langley (Eds.), Computational Models of Scientific Discovery and Theory Formation. San Mateo, CA: Morgan Kaufmann Publishers. Also appears in Sharlik & Dietterich (Eds.), Readings in Machine Learning. San Mateo, CA: Morgan Kaufmann Publishers.
25. Falkenhainer, B. (1993). Ideal Physical Systems. Proceedings of AAAI-93. 600-605.
26. Falkenhainer, B., & Forbus, K. (1991). Compositional Modeling: Finding the Right Model for the Job. Artificial Intelligence, 51. 95-143.
27. Faltings, B. (1987). Qualitative Kinematics in Mechanisms. Proceedings of IJCAI-87. 436-443.
28. Faltings, B., & Struss, P. (Eds.) (1992). Recent Advances in Qualitative Physics. Cambridge, Mass.: MIT Press.
29. Faltings, B., & Sun, K. (1993). Computer-Aided Creative Mechanism Design. Proceedings of IJCAI-93. 1451-1457.
30. Forbus, K. (1980). Spatial and Qualitative Aspects of Reasoning and Motion. Proceedings of AAAI-80. 170-173.
31. Forbus, K. (1984). Qualitative process theory. Artificial Intelligence, 24. 85-168.
32. Forbus, K. (1989). Introducing Actions into Qualitative Simulation. Proceedings of IJCAI-89. 1273-1278.
33. Forbus, K., & Falkenhainer, B. (1990). Self Explanatory Simulations: An integration of qualitative and quantitative knowledge. Proceedings of AAAI-90. 380-387.
34. Forbus, K., & Falkenhainer, B. (1995). Scaling up Self-Explanatory Simulators: Polynomial-time compilation. Proceedings of IJCAI-95. 1798-1805.
35. Forbus, K., & Gentner, D. (1986). Learning physical domains: Towards a theoretical Framework. In R. Michalski, J. Carbonell, & T. Mitchell (Eds.), Machine Learning: An Artificial Intelligence Approach, Volume 2. Morgan-Kaufmann. 311-348.
36. Forbus, K., Nielsen, P., & Faltings, B. (1991). Qualitative spatial reasoning: The CLOCK project. Artificial Intelligence, 51, 417-471.
37. Forbus, K., & Whalley, P. (1994). Using Qualitative Physics to Build Articulate Software for Thermodynamics Education. Proceedings of AAAI-94. 1175-1182.
38. Frederiksen, J., & White, B. (1990). Causal Model Progressions as a Foundation for Intelligent Learning Environments. Artificial Intelligence, 42. 99-157.
39. Funt, B. (1980). Problem-Solving with Diagrammatic Representations. In J. Glasgow, B Karan, & N. Narayanan (Eds.), Diagrammatic Reasoning. 33-68. Cambridge, Mass.: AAAI Press/ The MIT Press.
40. Gentner, D., & Stevens, A. (Eds.). (1983). Mental Models. Hillsdale, NJ: Erldaum.
41. Glasgow, J., Karan, B., & Narayanan, N. (Eds.). (1995). Diagrammatic Reasoning. Cambridge Mass.: AAAI Press/The MIT Press.
42. Glasgow, J., & Papadias, D. (1995). In J. Glasgow, B. Karan, & N. Narayanan (Eds.), Diagrammatic Reasoning. Cambridge, Mass.: AAAI Press/ The MIT Press.

43. Guerrin, F. (1991) Qualitative reasoning about an ecological process: Interpretation in hydroecology. Ecological Modeling 59: 165-201.
44. Hogge, J.C. (1987). Compiling Plan Operators from Domains Expressed in Qualitative Process Theory. Proceedings of IJCAI-87 229-233.
45. Hayes, P. (1979). The naive physics manifesto. In D. Michie (Ed.) Expert Systems in the Micro-Electronic Age Edinburgh University Press.
46. Hayes, P. (1985). Naive Physics 1: Ontology for liquids. In Hobbs, R., & Moore, R. (Eds.), Formal Theories of the Commonsense World Norwood, NJ: Ablex Publishing Corporation.
47. Hollan, J., Hutchins, E., & Weitzman, L. (1984). STEAMER: an interactive inspectable simulation-based training system. AI Magazine, Summer 1994.
48. Iwasaki, Y., & Simon, H. (1986). Theories of Causal Observing: Reply to de Kleer and Brown. Artificial Intelligence, 29(1). 63-68.
49. Iwasaki, Y. & Low, C.M. (1992). Device Modeling Environment: An Integrated Model Formulation and Simulation Environment for Continuous and Discreet Phenomena. Proceedings of CISE-92
50. Iwasaki, Y., Farquhar, A., Saraswat, V., Bobrow, D., & Gupta, V. (1995). Modeling Time in Hybrid Systems: How Fast is Instantaneous"? Proceedings of IJCAI-95
51. Iwasaki, Y., Tessler, S., & Law, K. (1995). Qualitative Structural Analysis Through Mixed Diagrammatic and Symbolic Reasoning. In J. Glasgow, B. Karan, & N. Narayanan (Eds.), Diagrammatic Reasoning Cambridge, Mass.: AAAI Press/ The MIT Press. 711-729.
52. Joscowicz, L. (1990). Mechanism Comparison and Classification for Design. Research in Engineering Design, 1(2).
53. Joscowicz, L., & Sacks, E. (1993). Automated Modeling and Kinematic Simulation of Mechanisms. Computer Aided Design, 25(2).
54. Joscowicz, L., & Sacks, E. (1993). Configuration Space Computation for Mechanism Design. IEEE-94.
55. Kamps, J., & Peli, G. (1995). Qualitative Reasoning beyond the Physics Domain: The Density Dependence Theory of Organizational Ecology. Proceedings of QR-95 114-121.
56. Kay, H., & Kuipers, B. (1993). Numerical Behavior Envelopes for Qualitative Models. Proceedings of AAAI-93. 606-613.
57. Kim, H. (1990). Qualitative Reasoning about Geometry of Fluid Flow. Proceedings of Cognitive Science Society-90 117-124.
58. Kim, H. (1992). Extending the Contained-Stuff Ontology with Geometry. In Proceedings of 6th International Workshop on Qualitative Reasoning. Edinburgh, Scotland.
59. Kim, H. (1992). Qualitative Kinematics of Linkages. In B. Faltings & P. Struss (Eds.), Recent Advances in Qualitative Physics Cambridge, Mass.: The MIT Press.
60. Kuipers, B. (1986). Qualitative Simulation. Artificial Intelligence, 29. 289-338.
61. Kuipers, B. (1994). Qualitative Reasoning: Modeling and simulation with incomplete knowledge. Cambridge, Mass.: MIT Press.
62. Kuipers, B. & Berleant, D. (1988). Using Incomplete Quantitative Knowledge in Qualitative Reasoning. Proceedings of AAAI-88 324-329.
63. Kuipers, B., & Byun, Y. (1991). A Robot Exploration and Mapping Strategy Based on Semantic Hierarchy of Spatial Reasoning. Journal of Robotics and Autonomous Systems, 8. 47-63.

64. Lang, K.R. Moore, J.C., and Whinston, A.B. 1995. Computational systems for qualitative economics. Computational Economics 8: 1-26.
65. Le Clair, S., Abrams, F., & Matejka, R. (1989). Qualitative Process Automation: Self Directed Manufacture of Composite Materials. AI EDAM, 3(2). 125-136.
66. Mavrovouniotis, M., & Stephanopoulos, G. (1988). Formal order-of-magnitude reasoning in process engineering. Computers and Chemical Engineering, 12 (9/10). 867-881.
67. Milne, R., & Trave-Massuyes, L. (1993). Application oriented Qualitative Reasoning. The Seventh International Workshop on Qualitative Reasoning about Physical Systems. 145-156.
68. Nayak, P. (1993). Order of Magnitude Reasoning Using Logarithms. In Proceedings of 7th Int. Workshop on Qualitative Reasoning, Orcas Island, WA.
69. Nayak, P. (1994). Causal Approximations. Artificial Intelligence, 70. 277-334.
70. Nielsen, P. (1988). A qualitative approach to mechanical constraint. Proceedings of AAAI-88.
71. Pisan, Y. (1995). A visual routines based model of graph understanding. In Proceedings of the Seventeenth Annual conference of the Cognitive Science Society (pp. 692-697). Hillsdale, NJ: Lawrence Erlbaum Associates.
72. Price, C., Pugh, D. Wilson, M. and Snooke, N. The FLAME system: Automating electrical failure modes and effects analysis (FMEA). Proc. Annual Reliability and Maintainability Symposium, pp 90-95, IEEE 1995.
73. Raiman, O. (1991). Order of magnitude reasoning. Artificial Intelligence, 51. 11-38.
74. Reece, S., & Durrant-Whyte, H. (1995). A Qualitative Approach to Sensor Data Fusion for Mobile Robot Navigation. Proceedings of IJCAI-95. 36-41.
75. Rickel J., & Porter, B. (1994). Automated Modeling for Answering Prediction Questions: Selecting the Time Scale and System Boundary. Proceedings of AAAI-94. 1191-1198.
76. Salles, P., Muetzelfeldt, R., & Pain, H. (1995). Qualitative Models in Ecology and their use in Intelligent Tutoring Systems. Presented at the workshop, "The Use of Qualitative Reasoning Techniques in Interactive Learning Environments" at AIED-95.
77. Sgorous, N. (1993). Representing Physical and Design Knowledge in Innovative Engineering Design (Tech. Rep. No. 49). Evanston: Northwestern University, Institute for the Learning Sciences.
78. Shen, Q., & Leitch, R. (1993). Fuzzy Qualitative Simulation. IEEE Trans. on Systems, Man, and Cybernetics, 23(4). 1038-1061.
79. Shimomura, Y., Tanigawa, S., Umeda, Y., & Tomiyama, T. (1995). Development of Self-Maintenance Photocopiers. Proceedings of IAAI-95. 171-180.
80. Sime, J., & Leitch, R. (1992). A Learning Environment Based on Multiple Qualitative Models. Presented at Intellectual Tutoring Systems Conference.
81. Skorstad, G. (1992). Towards a Qualitative Lagrangian Theory of Fluid Flow. In Proceedings of AAAI-92. 691-696.
82. Struss, P. (1988). Mathematical Aspects of Qualitative Reasoning. Int. J. Artificial Intelligence in Engineering, 3(3).156-169.
83. Umeda, Y., Tomiyama, T., & Yoshikawa, H. (1992). A Design Methodology for a Self-Maintenance Machine Based on Functional Redundancy. First International Conference on Intelligent Systems Engineering. Conference Publication No. 360. 179-184.

84. Vadillo, J.A., Diaz-Illarraza, A., Fernandez-Gastro, I., Gutierrez, J., and Elorriaga, J.A. 1994. The role of domain knowledge representation for explanation in ITTS. Proceedings of CALISCE'94. Paris, 245-252.
85. Weld, D. (1990). Theories of Comparative Analysis. Cambridge, Mass.: MIT Press.
86. Weld, D., & de Kleer, J. (Eds.). (1990). Readings in Qualitative Reasoning about Physical Systems. Los Altos, CA: Morgan Kaufmann.
87. Williams, B. (1991). A Theory of Interactions: Unifying Qualitative and Quantitative Algebraic Reasoning. Artificial Intelligence. 39-94.
88. Yannou, B., & Vasiliu, A. (1995). Design Platform for Planar Mechanisms based on a Qualitative Kinematics. Proceedings of QR-95. 191-200.
89. Yip, K. (1991). KAM: A System for Intelligently Guiding Numerical Experimentation by Computer. Artificial Intelligence Series. Cambridge, Mass.: MIT Press.
90. Yip, K. (1995). Reasoning about Fluid Motion I: Finding Structures. Proceedings of IJCAI-95. 1782-1788.
91. Zhao, F. (1994). Extracting and Representing Qualitative Behaviors of Complex Systems in Phase Space. Artificial Intelligence, 69. 51-92.

8. Further Information

A good introduction to qualitative physics is [Weld & de Kleer, 1990], which provides access to many of the classic papers in the field. [Faltings & Struss, 1992] provides a sample of more recent papers. An excellent textbook on the QSIM approach to qualitative physics is [Kuipers, 1994]. For an introduction to diagrammatic reasoning, see [Glasgow et al., 1995]. [Milne & Trave, 1993] provides an extensive survey of application-oriented qualitative reasoning work.

Papers on qualitative reasoning routinely appear in Artificial Intelligence, AI in Engineering Design and Manufacturing (AIEDAM), and IEEE Expert. Many papers first appear in the proceedings of AAAI, IJCAI, ECAI. Every year there is an International Qualitative Reasoning Workshop, whose proceedings document the latest developments in the area. Proceedings for a particular workshop are available from its organizers.