AAAI-88
The Seventh National Conference on Artificial Intelligence

Using Incomplete Quantitative Knowledge in Qualitative Reasoning*

Benjamin Kuipers and Daniel Berleant
Department of Computer Sciences
University of Texas, Austin, Texas 78712
email: kuipers@cs.utexas.edu *or* berleant@sally.utexas.edu

Abstract

Incomplete knowledge of the structure of mechanisms is an important fact of life in reasoning, commonsense or expert, about the physical world. Qualitative simulation captures an important kind of incomplete, ordinal, knowledge, and predicts the set of qualitatively possible behaviors of a mechanism, given a qualitative description of its structure and initial state. However, one frequently has *quantitative* knowledge as well as qualitative, though seldom enough to specify a numerical simulation.

We present a method for incrementally exploiting incomplete quantitative knowledge, by using it to refine the predictions of a qualitative reasoner. Incomplete quantitative descriptions (currently ranges within which unknown values are assumed to lie) are asserted about some landmark values in the quantity spaces of qualitative parameters. Unknown monotonic function constraints may be bounded by numerically computable envelope functions. Implications are derived by local propagation across the constraints in the model.

When this refinement process produces a contradiction, a qualitatively plausible behavior is shown to conflict with the quantitative knowledge. When all predicted behaviors of a given model are contradicted, the model is refuted. If a behavior is not refuted, propagation of quantitative information results in a mixed quantitative/qualitative description of behavior that can be compared with other surviving predictions for differential diagnosis.

1 Introduction

A qualitative model of a device or system is an abstraction of a set of real systems. The behavior of these systems can vary greatly, yet purely qualitative descriptions of these behaviors are identical. Quantitative knowledge about these systems can, however, allow them and their behaviors to be distinguished. Adding quantitative information to qualitative modeling allows more precise characterization of systems and their behaviors. This increased precision can help in diagnosis and prediction of behavior, even if there is much too little for numerical simulation. Consider the example of the pressure regulator.

A working pressure regulator of the type described in [de Kleer and Brown, 1985] has a fluid input of varying pressure. The regulator has an internal, continuously variable valve which regulates the pressure at the output port so that its variation is considerably smaller than the variation at the input port. It does this by sensing the output pressure and closing the valve to restrict the flow as the output pressure rises, or opening the valve when the output pressure falls.

The pressure regulator may malfunction by having the internal valve stuck in one position, so that it cannot control the output pressure. For both the working and stuck pressure regulators, an increase in input pressure leads to an increase in output pressure, but for the working regulator the increase is significantly less. Qualitative simulation of both the working and stuck regulators indicates correctly that an input pressure increase implies an output pressure increase. However, the qualitative description alone cannot distinguish between the working and stuck regulators on the basis of output pressure variation (or any other easily measured parameter). This problem may arise with any proportionally controlled, negative feedback system, an important class of real mechanisms.

Quantitative information is necessary to resolve this ambiguity, but we wish to preserve our ability to reason reliably with incomplete knowledge of the structure and numerical values characterizing the physical system. Our mixed qualitative-quantitative reasoner, Q2, makes it possible to assert incomplete quantitative knowledge in the form of ranges[1], about the landmark values in Kuipers' [1986] QSIM behavioral description, and propagate their consequences. Our method is applicable to other qualitative reasoning systems with limitations discussed in section 4.

In the case of the pressure regulator, we assumed plausible ranges[2] of values for resistance and flow capacity, and

*This research was supported in part by the National Science Foundation through grant number IRI-8602665.

[1]Fully specified quantitative values are expressed as ranges whose endpoints are identical.

[2]The term *range* is used rather than *interval*, because the rules of interval arithmetic are not always valid in this application. In interval arithmetic, if XY=K for intervals X, Y and K, the width of Y *decreases* for increased width of X given K (cf. Alefeld & Herzberger [1983]). But if X and Y are ranges representing reals whose values are uncertain, then increased width for X represents greater uncertainty in X, hence greater uncertainty (i.e. *increased* width) in Y. Our ranges are to be interpreted as representing probability distribution functions. Range $[A, B]$ thus represents any pdf whose value is positive from A to B, and zero otherwise.

simulated the response of the regulator to a doubling of the input pressure from $[5, 5.1]$ to $[10, 10.2]$. Each of the two models (working and stuck) predicted a single qualitative behavior: Output pressure increased. Augmenting the qualitative descriptions with quantitative ranges, the working model predicted the final value of the output pressure to be in $[1.91, 2.98]$, while the stuck model predicted an output pressure in $[3.8, 6.2]$. This is precisely what is required for differential diagnosis between the two models.

2 Propagation of Incomplete Quantitative Information

We will explain our quantitative propagation method in the context of a simple one-tank "bathtub" system; in this case one with a partially blocked drain, so that outflow increases only slowly with pressure.

There are three distinct qualitative behaviors for a bathtub which is being filled from empty with the drain left open: (1) equilibrium between inflow and outflow before amount reaches FULL, (2) overflow while inflow is greater than outflow, and (3) equilibrium between inflow and outflow exactly when amount reaches FULL.

In Q2, two types of quantitative information are provided as part of the initial description of the system:

- Quantitative ranges describing what is known about the values of certain landmark values, in this case the landmark IF* of the parameter $inflow(t)$, and the landmark TOP of the parameter $level(t)$.
- Numerically computable envelopes that bound the (unknown and possibly nonlinear) monotonic function constraints, such as $outflow = M^+(pressure)$.

Figure 1 shows the only quantitatively consistent behavior out of the three qualitative possibilities, given initial quantitative assertions about TOP, IF*, and envelopes constraining the relations between amount and level, level and pressure, and pressure and outflow. The two equilibrium behaviors were found to be inconsistent with the quantitative information given.

2.1 Types of Quantitative Propagation

Quantitative propagation occurs in different ways for the various qualitative constraints being propagated over. As a notational convention, if the qualitative behavior has $parameter(t) = L$ for a landmark L at a particular time-point t, we may use either $parameter(t) = [lo, hi]$ or $L = [lo, hi]$, to indicate that the quantitative range $[lo, hi]$ must contain the (unknown) numerical value of L.

In Q2, each type of qualitative constraint is associated with a procedure for propagating partial quantitative information among its arguments. These procedures define a quantitative semantics for the constraint that must of course be consistent with the semantics already defined by the qualitative simulator. The four types of methods for propagating incomplete quantitative information are:

1. **Propagation across arithmetic constraints: ADD, MULT, MINUS.**
   This is exemplified by an ADD constraint in a model of a bathtub, as shown in table 1. Note that divide and (binary) subtract constraints are trivially implemented with ADD and MULT.

(a) An ADD constraint:
netflow = inflow − outflow
(b) Landmark values at time t_1 (see Fig. 1):
netflow(t_1) = inflow(t_1) − outflow(t_1)
i.e., NF-1 = IF − OF-1*
(c) in terms of known ranges:
[0.051, 0.146] = [1, 1.01] − [0, 9999]
(d) The ADD can narrow the range for outflow(t_1):
[0.051, 0.146] = [1, 1.01] − [0.864, 0.948]

Table 1: Propagation reduces range bounds for each landmark.

2. **Propagation across monotonic function constraints: M^+, M^-.**
   This is typified in the bathtub model by an "M+" monotonic constraint between *amount* of water and *level* in the tub, indicating that a change in either parameter implies a change in the other in the same direction. A qualitative monotonic function is a generalization of a large space of possible quantitative functions - indeed, all monotonic quantitative functions for which the monotonicity has the same sign as that of the corresponding qualitative function. There is a middle ground between purely qualitative and fully specified quantitative monotonic functions. We implement this middle ground by using upper and lower ENVELOPES (figure 2).
   ENVELOPES are quantitative functions which bound the space of quantitative functions that could apply to a monotonic constraint to a greater extent than the sign of the monotonicity. For the bathtub example, a particular tub may be consistent with a bathtub model that is partly quantified by envelopes constraining the relation between amount and level if its function relating amount and level falls within those envelopes. Otherwise it is definitely not consistent (maybe it is very funny-shaped tub, or perhaps not a tub at all but a sink or swimming pool). Propagation through a partially quantified M+ constraint occurs as described in figure 2.
3. **Propagation across quantity spaces.**
   Consider the $netflow$ of water into the tub. At time T1, the value of $netflow$ is whatever quantitative value is associated with the landmark named "NF-1" (fig. 1). This value must be less than the value of NF-0, which may be as high as 1.01, but is greater than 0. Thus from the ordinal position of NF-1 and the quantitative information associated with its neighbors, we infer that $netflow = [0, 1.01]$. With the help of other sources of constraint, propagation eventually narrows it all the way to $[0.051, 0.146]$.
4. **Propagation across time-points: D/DT.**
   Finally, there is information flow from one state in the behavior of a model to another. This occurs via D/DT constraints, e.g., D/DT(amount)=netflow. By looking at quantitative information about the values

Structure: Standard bathtub with open drain.,
Initialization: Filling at constant rate, starting from empty (S-0)
Behavior 2 of 3. Final state: (TRANSITION), (TRANSITION-IDENTITY NIL), T<INF.
Time point ranges: T0=[0.0 0.0] T1=[0.873 18.216]

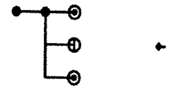

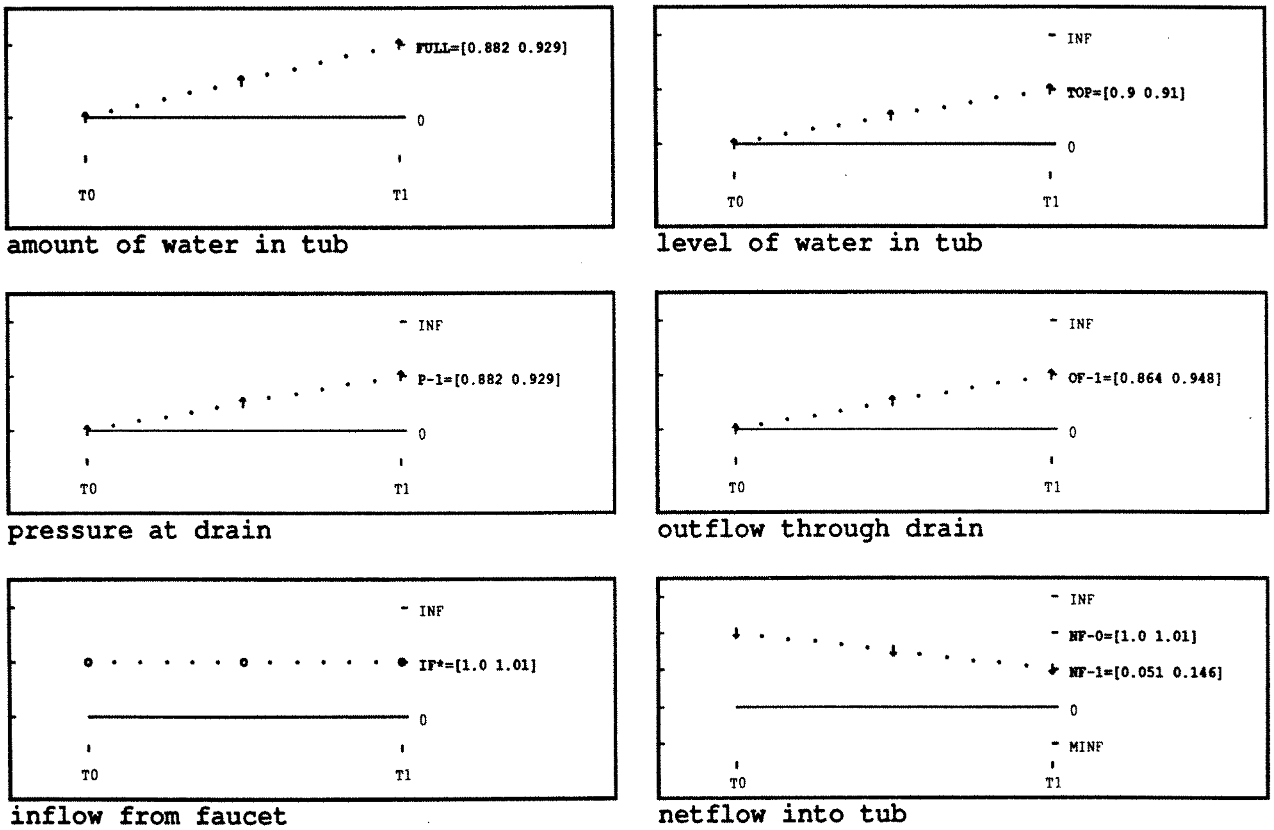

Figure 1

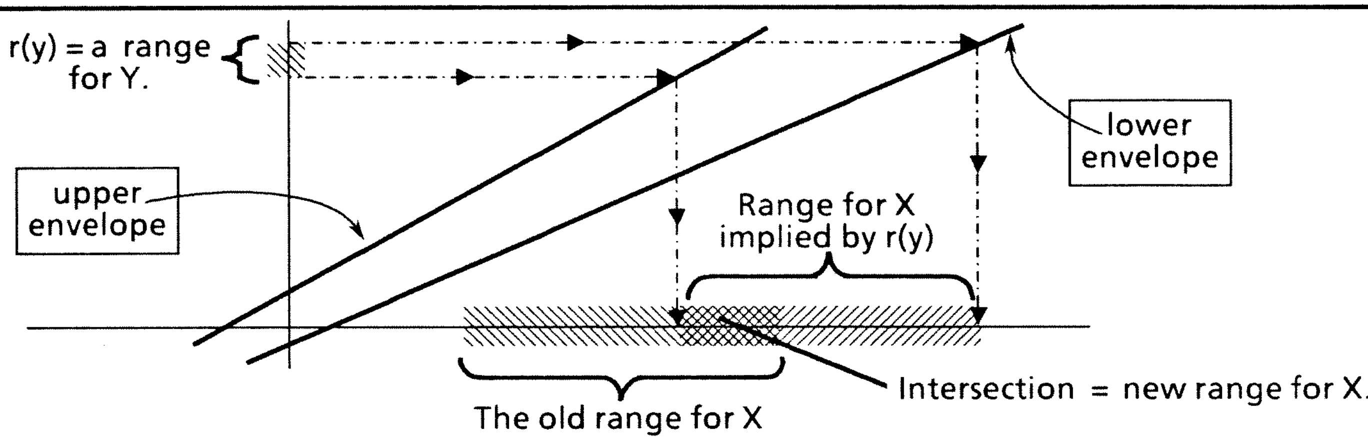

Figure 2: Range propagation across an M^+ constraint: $Y = M^+(X)$.

Given r(y), a range for Y, find the new range for X:

1) Project r(y) across the envelopes.
2) Intersect the projection with the old range for x.

at adjacent time points of the integral, the derivative and the time, propagation can potentially constrain the ranges associated with each of these. For the bathtub, the mean value theorem of calculus tells us that $\exists T^* \in (T0, T1)$ such that

$$netflow(T^*) = \frac{amount(T1) - amount(T0)}{(T1 - T0)}$$

From figure 1 we see that $netflow(T^*) = [0.051, 1.01]$. We also see that amount started out at 0 and climbed to anywhere from 0.882 to 0.929. T0 is known to have the value 0. Thus,

$$T1 = 0 + \frac{[0.882, 0.929]}{[0.051, 1.01]} = [.873, 18.216]$$

These four kinds of constraint apply the quantitative information provided by the user to narrow the ranges associated with each landmark of each parameter until either no further narrowing is possible, or an inconsistency is flagged. Inconsistency, of course, is relative to a behavior and means that the behavior is not compatible with the available quantitative information. If all behaviors of a model are inconsistent then an additional inference is possible: The model itself is incompatible with the quantitative information, whether that information is known *a priori* or from observations.

2.2 The propagation algorithm

The range propagator (cf. [Davis, 1987]) is straightforward, making no distinction between the various kinds of constraint for control purposes. It starts by setting each landmark of each model parameter to an initial range of $[0+, \infty]$, $[-\infty, 0]$, or $[0, 0]$, depending on whether the landmark is above zero, is the "0" landmark, or is below zero. Then any quantitative information provided by the user is used to narrow the appropriate landmarks. For the bathtub, (inflow . IF*), the "IF*" landmark of the inflow of water from the faucet, is initialized to $[1.0, 1.01]$. In addition, (level . TOP) is initialized to $[0.9, 0.91]$, meaning that we are dealing with bathtubs whose height falls between 0.9 and 0.91.

Narrowed landmarks can potentially enable narrowing of other landmarks. A constraint is *attached* to a landmark L if it and range r(L) can be used to try to narrow other landmarks. All constraints attached to the narrowed landmarks are added to an agenda. The propagation algorithm now takes the first constraint off the agenda and uses it to try to narrow the landmarks associated with it. If it fails it goes back to the agenda for the next constraint. If it succeeds it adds to the agenda all constraints attached to any landmarks it succeeded in narrowing, and returns to the agenda for a new constraint to process. The current implementation is depth first, and termination occurs when the agenda is empty. Our models run in on the order of 1 minute.

2.3 Reasoning with Models, Behaviors, and Values.

The Q2 reasoner can make distinctions at three levels of granularity. The coarsest level deals with different systems and their MODELS. For the bathtub system, different models might include bathtubs with completely blocked drains, ones with rusted out bottoms, ones with partially blocked drains, and upside-down bathtubs.

Qualitative-quantitative reasoning can also discriminate among different qualitative BEHAVIORS of a model. Behaviors of a model are consistent or not with the incomplete quantitative knowledge a user has provided. We have previously shown how Q2 can infer, for a bathtub with partially blocked drain that satisfies the specifications of certain quantitative ranges, that the equilibrium behaviors are inconsistent and only the overflow behavior could occur.

Figure 3 illustrates the reasoning about models and behaviors that is one of the capabilities of Q2.

The finest level of granularity deals with VALUES of parameters. For the bathtub, quantitative data in the form of envelopes, and ranges that constrain uncertainty about the tub height and the faucet flow, imply quantitative predictions (figure 1) about other qualitative aspects of the bathtub and its behavior.

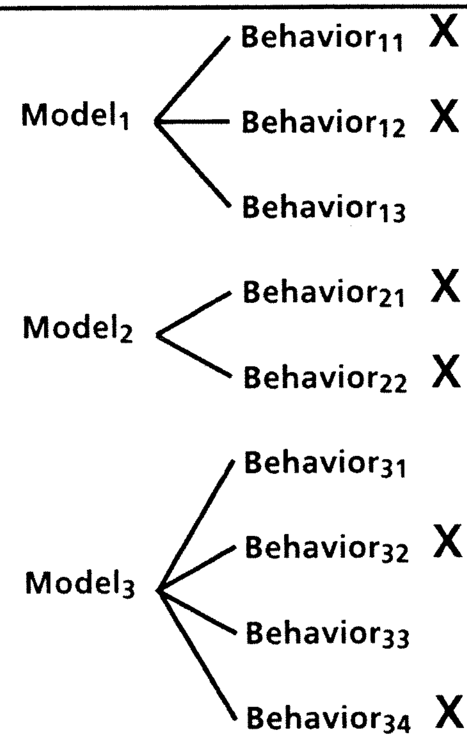

Figure 3:
Quantitative reasoning eliminates behaviors, and sometimes an entire model.
X designates an inconsistent behavior.

3 A More Complex Example

Consider the more complex example of a one-tank equilibrium system with a proportional controller attempting to keep *amount* near a desired point by modifying *inflow*. One example of such a system is a heating system where the rate of heat inflow from the heater is proportional to the difference between the actual and desired temperature of the heated vessel. (The usual household thermostat does on-off control, not proportional control.) Linear proportionality is only a special case of such a controller: In general, restoring force may be a monotonic function of the measured error. Proportionally controlled systems are very common in the world, including physiological mechanisms, chemical systems, automobile cruise control, etc.

We created four distinct models for a hypothetical proportionally controlled heating system:

1. The properly working system;
2. Continuous maximum heating, regardless of temperature;
3. No heat at all, regardless of temperature;
4. Thermostat with faulty calibration, which acts as though the temperature is higher or lower than it really is and therefore causes an equilibrium temperature different from the thermostat setting.

Most models have more than one possible qualitative behavior. For example, a properly working temperature controlling system may respond successfully to a demand for increased heating, or it may "max out" by delivering heat steadily at its maximum capacity despite increasing demand. When given a particular set of *a priori* and observed quantitative knowledge, Q2 generated eighteen qualitatively possible behaviors from the four models, and used the quantitative knowledge to eliminate all but two of them. The remaining two make identical predictions, since the fault model accounting for the uncalibrated thermostat includes the behaviors of the properly working thermostat as special cases.

4 Related Work

There has been considerable other work relevant to the integration of quantitative with qualitative knowledge.

The measurement interpretation methods developed by Forbus [1983, 1986] are closest to our work in terms of the problem solved, though quite different in approach. We, like Forbus, are attempting to interpret quantitative measurements by matching the observed measurements against the predictions of a model. Where there are several candidate models, or several behaviors of a given model, failure to match refines the set of remaining viable candidates.

Our method differs from Forbus' approach in the handling of quantitative information. In the more complete formulation [Forbus, 1986], a continuous stream of quantitative data is mapped into a stream of qualitative descriptions; in his example, directions of change, or Ds values $\{+1, 0, -1\}$. In an example involving heating a container of mixed alcohol and water, the stream of temperature measurements is described qualitatively as $[+1, 0, +1, 0, +1]$. The total envisionment of a given situation can be regarded as a finite-state transition graph, which is used to "parse" the stream of Ds values from an acceptable initial state to an acceptable final state. The path successfully taken through the envisionment describes the sequence of process structures the system goes through. Failure to parse presumably refutes the model. Notice that a significant amount of quantitative information is lost when fine-grained quantitative measurements are mapped to coarser-grained qualitative representations (in this case the Ds values), and the comparison with the model takes place with the measurements expressed in the same coarse qualitative terms as those used in the model.

In Q2, by contrast, quantitative information is used to *augment* the qualitative descriptions. One advantage of this approach is that quantitative information may be propagated across constraints, providing information about landmarks of parameters whose values have not even been measured. Another advantage is that uncertain quantitative knowledge can be expressed precisely (in the form of ranges), and used effectively. Third, multiple behaviors that satisfy the known quantitative constraints now carry quantitative predictions, easing the problem of differential diagnosis. Q2 is not currently designed to reason with a *continuous* stream of quantitative measurements. Instead, it takes descriptions of values at the endpoints of time-periods of monotonic change. However, utilizing measurements (as well as making partial quantitative predictions) for arbitrary times is a direction for future work, both due to its promise for practical application and because it appears to be a natural extension of the current system.

Karp and Friedland [1987] also share the goal of integrating qualitative and quantitative constraints in reasoning about mechanisms. They create a frame for each parameter at each instant, capable of representing a rich variety of algebraic equations and inequalities involving that value, plus frames for interactions between constraints. While the expressive and inferential power of their approach is potentially very large, so is the potential for combinatorial explosion, since there is no clear structure on the types of constraints and the circumstances under which different types of constraints are applied. In Q2, ordinal relations between values and landmarks are used by QSIM to propose qualitatively possible behaviors, and quantitative ranges are then used to refine or refute each behavior. The use of distinct types of knowledge for distinct purposes supports conceptual clarity and implementational efficiency.

Simmons' [1986] quantity lattice, and Sacks' [1987] hierarchical inequality reasoner are more powerful methods of arithmetic reasoning than the package currently in Q2. We plan extensions along these lines.

As discussed above, our method depends on starting with a qualitative description of behavior in terms of landmark values which function as "names" for real numbers, and about which we can accumulate and refine quantitative descriptions. Thus, our approach does not apply in any natural way to qualitative category representations such as {high, medium, low}, since these symbols refer to sets rather than values, and the boundaries between the sets are not distinctive values. Furthermore, qualitative category representations do not support a rigorous form of qualitative simulation, since limit analysis is not meaningful in that context. It is also relatively difficult to apply our approach to the de Kleer and Brown [1985] $\{+, 0, -\}$ representation, since the quantity space contains no non-zero landmarks, and zero already has a precise value.

5 Conclusions and Directions for Future Work

For conceptual clarity during development, the current implementation of Q2 applies quantitative knowledge to individual, complete qualitative behaviors from the output of QSIM. We plan to interleave quantitative and qualitative processing, so that quantitative inferences can be applied to partially complete qualitative behaviors. Where a quantitative inconsistency can be identified at an early stage, an entire subtree of qualitative behaviors may be eliminated, greatly increasing the efficiency of the overall simulation.

Our current implementation represents incomplete quantitative knowledge as numerically bounded ranges. We believe that our approach will also be applicable to propagation of quantities described by probability distributions (i.e. mean and variance) [J. Pearl and P. Cheeseman, personal communication]. In this case, the result will not be to filter out certain behaviors as inconsistent, but to define a probability distribution across the set of possible behaviors. Reasoning with mean-variance descriptions of quantities is of obvious practical importance, given the probabilistic nature of most real-world measurements.

As we have discussed, after assimilating a set of quantitative observations, the refined quantitative descriptions of surviving behaviors are precisely what is needed for differential diagnosis, for example by selecting a quantity whose ranges in two different behaviors are non-overlapping, and testing for its value. The work on "diagnosis from first principles" by Davis [1984], Genesereth [1984] and Reiter [1987] provides methods for optimizing the selection of new tests.

It should also be possible to perform a sensitivity analysis [Raiffa, 1970] on the results of the propagation, to assess the sensitivity of Q2's conclusions to variations in the quantitative observations. This will provide a first step towards capturing second-order uncertainty in the descriptions of incomplete quantitative knowledge.

6 References

[Alefeld and Herzberger, 1983] Gotz Alefeld and Jurgen Herzberger. *Introduction to Interval Computations.* Academic Press, 1983, p.5.

[Davis, 1984] Randall Davis. Diagnostic reasoning based on structure and behavior. *Artificial Intelligence*, 1984, **24**: pages 347-410.

[Davis, 1987] Ernest Davis. Constraint Propagation with Interval Labels. *Artificial Intelligence*, 1987, **32**: pages 281-331.

[de Kleer and Brown, 1985] Johan de Kleer and John Seely Brown. A Qualitative Physics Based on Confluences. In *Qualitative Reasoning about Physical Systems*, ed. by Daniel C. Bobrow. MIT Press, 1985.

[Forbus, 1983] Kenneth D. Forbus. Measurement Interpretation in Qualitative Process Theory. In *Proceedings IJCAI-8*, 1983, pages 315-320.

[Forbus, 1986] Kenneth D. Forbus. Interpreting Measurements of Physical Systems. In *Proceedings AAAI-86*, 1986, pages 113-117.

[Genesereth, 1984] Michael R. Genesereth. The use of design descriptions in automated diagnosis. In *Artificial Intelligence*, 1984, **24**: pages 411-436.

[Karp and Friedland, 1987] Peter D. Karp and Peter Friedland. Coordinating the Use of Qualitative and Quantitative Knowledge in Declarative Device Modeling. Knowledge Systems Lab. Report KSL 87-09, Stanford Computer Science Dept., 1987.

[Kuipers, 1986] Benjamin J. Kuipers. Qualitative Simulation. In *Artificial Intelligence*, 1986, **29**: pages 289-338.

[Raiffa] H. Raiffa. *Decision Analysis: Introductory Lectures on Choices under Uncertainty.* Addison-Wesley, Reading MA, 1970.

[Reiter, 1987] Raymond Reiter. A Theory of Diagnosis from First Principles. In *Artificial Intelligence*, 1987 **32**: pages 57-95.

[Sacks, 1987] Elisha Sacks. Hierarchical reasoning about inequalities. In *Proc. AAAI-87*, pages 649-654.

[Simmons, 1986] Reid G. Simmons. "Commonsense" Arithmetic Reasoning. In Proc. AAAI-86, pages 118-124.

[Widman] Lawrence E. Widman. Representation Method for Dynamic Causal Knowledge using Semi-Quantitative Simulation. In *Proceedings of the Fifth World Conference on Medical Informatics (MEDINFO86)*, pages 180-184.