Interpreting measurements of physical systems

Kenneth D. Forbus
Qualitative Reasoning Group
Department of Computer Science
University of Illinois
1304 W. Springfield Avenue
Urbana, Illinois, 61801

Abstract

An unsolved problem in qualitative physics is generating a qualitative understanding of how a physical system is behaving from raw data, especially numerical data taken across time, to reveal changing internal state. Yet providing this ability to "read gauges" is a critical step towards building the next generation of intelligent computer-aided engineering systems and allowing robots to work in unconstrained envirionments. This paper presents a theory to solve this problem. Importantly, the theory is domain independent and will work with any system of qualitative physics. It requires only a qualitative description of the domain capable of supporting envisioning and domain-specific techniques for providing an initial qualitative description of numerical measurements. The theory has been fully implemented, and an extended example using Qualitative Process theory is presented.

1. Introduction

Interpreting numerical data is an important part of monitoring, operating, analyzing, debugging, and designing complex physical systems. A person operating a nuclear power plant or propulsion plant must constantly read and interpret gauges to maintain an understanding of what is happening and take corrective action, if necessary. Designing a new system requires running numerical simulations (or building models of the system) and analyzing the results. Diagnosis requires interpreting behavior, both to see if the system is actually operating correctly and to determine if a hypothesized fault can account for the observed behavior. All of these problems require the ability to deduce the changing internal state of the system across time from measurements.

Currently there is a great deal of interest in applying qualitative physics to engineering tasks such as diagnosis (e.g., the articles in [Bobrow, 1985]). For such efforts to be successful, a theory about how to translate observed behavior, including numerical data, into useful qualitative terms is essential. This paper presents such a theory. The theory is domain independent and makes only two assumptions about the nature of the underlying domain model. Specifically, it assumes that:

1. Given a particular physical situation, a graph of all possible behaviors – an *envisionment* – may be generated.

2. Domain-specific criteria are available for quantizing numerical data into an initial qualitative description.

The theory is analogous to AI models of speech understanding (e.g., [Reddy, et. al, 1973]). In these models the speech signal is partitioned into segments, each of which is explained in terms of phonemes and words. Grammatical constraints are imposed between the hypothesized words to prune the possible interpretations. In this theory, the initial signal is partitioned into pieces which are interpreted as possible particular qualitative states of the system. By supplying information about state transitions, the envisionment plays the role of grammatical constraints, imposing compatibility conditions between the hypotheses for adjacent partitions.

1.1. Overview

The goal of this theory is to produce a general solution for the problem which can be instantiated for any particular physics and domain. Consequently we couch the analysis in an abstract vocabulary and specify what domain-dependent modules are required to produce initial qualitative descriptions. We demonstrate how the theory can be instantiated using an example involving *Qualitative Process theory* [Forbus, 1981, 1984]. The theory has been implemented, and the performance of the implementation on these examples is demonstrated. It should be noted that the theory and implementation have also been successfully applied to a completely different system of qualitative physics, the *qualitative state vector* ontology ([de Kleer, 1975], [Forbus, 1980, 1981b]), as described in [Forbus, 1986].

The next section provides a vocabulary for describing the initial data and places constraints on the segmentation process. Section 3 generalizes an earlier theory of interpreting measurements taken at an instant for QP theory [Forbus, 1983] and shows how the envisionment can be used to locally prune interpretations of segments. [1] Section 4 illustrates how global interpretations are constructed and how gaps in the input data can be filled. Finally we discuss planned extensions and some implications of the theory.

2. Input Data and Segmentation

First we describe the kinds of inputs the theory handles. We assume a function `time` which maps measurements to real numbers, and that the duration of an interval is simply the difference between the times for its start and end points. We also assume the temporal relationships described in [Allen, 1981] may be applied to intervals (i.e., `Meet`, `Starts`, and `Finishes`.) We say

`Observable(<p>, <i>)`

when property <p> can be observed in principle by instrument <i>. To say that we can measure the level of water in a can with our eyes we write[2]

`Observable(A[Level(C-S(water, liquid, can))], eyes)`

To say that some property is in fact observable at some time, we use the predicate `Observable-at`, which takes a time as an extra argument.

[1] An "envisionment" is, roughly, "the set of all qualitatively distinct possible behaviors of a system." However, sometimes it is used to refer to "all behaviors possible from some given initial state" and sometimes to "all behaviors inherent in some fixed collection of objects in some configuration, for each possible initial state." We call the first type *attainable* envisionments, and the second *total* envisionments. Here we are only be concerned with total envisionments.

[2] The first argument uses notation from QP theory; A is a function that maps from a quantity to a number representing the value of that quantity, Level is a function mapping from individuals to quantities, and C-S is a function denoting an individual composed of a particular substance in a particular state, distinguished by virtue of being in a particular place.

We say

Measured(<p>, <v>, <t>, <i>)

to mean that property <p> takes on value <v> at time <t>, as measured by instrument <i>. To say that we measured the level in can to be 5 centimeters 6 seconds after an experiment started, we write, ignoring units,

Measured(A[Level(C-S(water, liquid, can))],5,6,eyes)

We use these same conventions to define observations of arbitrary facts by interpreting them as functions whose range is {true, false}.

The input of a measurement interpretation problem is a set of *measurement sequences*, each consisting of a set of measurements totally ordered by the times of the measurements. Suppose we have some "grain" on time, st, such that events of duration shorter than st will not be considered relevant. (The problem of instantaneous events will be discussed in section 4.1.) Then two types of measurement sequences must be considered:

Close: The data is complete, in the sense that over the total interval of interest measurements are separated by durations no larger than st.

Sampled: There are temporal gaps in the data whose duration is larger than st.

Given an assumption of a finite "grain size" of analysis, with close data we are justified in assuming that contiguous segments of the data correspond to successive states of the system. With sampled data we can only make this assumption on close subsequences. *Regular* sequences are a subclass of close sequences where successive measurements are exactly st apart.

2.1. Segmenting the input data

The first problem is to partition the measurement sequences into meaningful pieces. We define a *segment* of a measurement sequence to be the largest contiguous interval over which the measured property is "constant". A symbolic property is constant over an interval if its value is identical for all measurements within that interval. Notice that in QP theory signs of derivatives are symbolic properties in this sense.

A numerical parameter is constant over a segment if the same *qualitative value* can be used to describe each measurement in the segment. The exact notion of qualitative value depends of course on the choice of domain representation and ontology.[3] All we require is that algorithms exist for taking numerical values and producing at least some qualitative description sanctioned by the representation used. In QP theory, for example, numerical values can be described in terms of inequalities, the *quantity space* representation. If some domain-specific constants are unknown, such as the boiling temperature of a particular substance at a certain pressure, partial information can be delivered. In the worst case, the sign of the derivative can be estimated.

Once numerical parameters are translated to qualitative values segmentation becomes simple. However, these segments cannot necessarily be identified with a single qualitative state. First, the qualitative value may be partial, as noted above. Second, a state transition may leave the measured parameters constant for some time (possibly forever). Consider a home heating system. Suppose you turn the thermostat up past the ambient temperature. If you cannot hear the furnace firing or touch a radiator, then you will not know for some time whether or not the system is actually working. This *hidden transition problem* must be taken into account when pruning interpretations, discussed below.

Each segment of a measurement sequence covers a non-empty collection of data, and since the data is temporally well-ordered there will be a maximum and minimum time associated with this data set. Let the minimum time be the start time and the maximum time be the end time. We define two intervals I_1, I_2 associated with distinct segments of a measurement sequence to be *adjacent* if there is no interval in between them (by assumption, I_1 and I_2 cannot overlap). If the minimum distance between the times of the end points is not greater than st, we also say that the intervals Meet, as defined in Allen's time logic. The function Int maps segments to intervals.

The local information provided by the segmentation of measurement sequences must be combined to form *global segments*, intervals over which the qualitative state of the system is not obviously different. We define global segments as follows. Let $\{MS_i\}$ be a collection of measurement sequences, each of which has a segmentation $\{S_{i,j}\}$. The global segmentation consists of a set of global segments $\{GS_k\}$ such that

1. The value of the property measured for each MS_i is constant over GS_k.

2. Starts(GS_k, Int($S_{i,j}$)) for some $S_{i,j}$, i.e., the start time of each global segment corresponds to the starting time of some segment in one or more of the segmented measurement sequences.

3. Finishes(GS_k, Int($S_{i,k}$)) for some $S_{i,k}$, i.e., the end time of each global segment corresponds to the end time of some segment in one or more of the segmented measurement sequences.

The first constraint prevents a global segment from straddling an obvious qualitative boundary, and the last two constraints ensure it spans the largest possible interval where qualitative values are constant. Thus global segments are good candidates for explanation by a single qualitative state.

2.2. QP Example, Part 1

Many changes in the physical world can be characterized as the result of physical processes, such as heat flow, liquid flow, boiling, and motion. Qualitative Process theory formalizes this intuitive notion of physical process and provides a qualitative language for differential equations that preserves distinctions required for causal reasoning.

QP theory provides several types of measurable properties, including the truth of predicates and relations, whether or not different processes are acting, and of course information about numbers. Ideally measurements of amounts and magnitudes should be segmented whenever their descriptions in terms of quantity spaces change. However, as we will see a great deal of information can be gleaned from just the signs of derivatives (i.e., the Ds value of a quantity, which ranges over {-1, 0, 1}).

Suppose we have a beaker that has a built-in thermometer. Suppose we also know that the beaker either contains some water, some alcohol, or a mixture of both. In this case we can always measure the temperature, i.e.

$\forall$ t $\in$ *time*

Observable-at(A[Temperature(inside(beaker))],
            thermometer, t)

If we plot the temperature with respect to time we might get the graph shown in Figure 1.[4]

If we don't know the numerical values for the boiling points of water and alcohol, then all we can get from this graph is the Ds value for temperature as a function of time. Providing this list of Ds values to the program results in six segments. Since this is the only property measured, each segment gives rise to a single global segment. The program's output is shown in Figure 2.

---

[3] To be a qualitative representation some such notion must exist; the primary purpose of such representations is to provide quantizations of the continuous world which form useful vocabularies for symbolic reasoning.

Figure 1 – **Temperature plotted as a function of time**

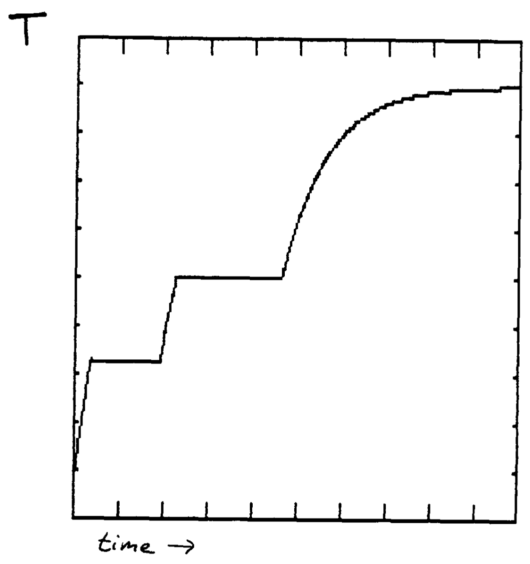

Figure 2 – **Segments and global segments for the QP problem**
Here are the segments and global segments generated by the implementation from the data in Figure 1.

```
ATMI: Finding initial segments...
  1 properties have been measured.
For Ds of (T INSIDE-BEAKER):
   Start time = 0.0, End time = 11.7.
   117 samples, taken 0.1 time units apart.
   Divided into 6 segments.
      Ds of (T INSIDE-BEAKER) is 1 from 0.0 to 1.3.
      Ds of (T INSIDE-BEAKER) is 0 from 1.4 to 2.1.
      Ds of (T INSIDE-BEAKER) is 1 from 2.2 to 4.1.
      Ds of (T INSIDE-BEAKER) is 0 from 4.2 to 5.8.
      Ds of (T INSIDE-BEAKER) is 1 from 5.9 to 8.5.
      Ds of (T INSIDE-BEAKER) is 0 from 8.6 to 11.7.
ATMI: Finding global segments...
There are 6 global segments.
```

---

3. Interpreting segments

If the segmentation based on domain-specific constraints is correct, a global segment should typically be explained as the manifestation of a single qualitative state. A qualitative state consists of a finite number of components,[5] some fraction of which are fixed by the measurement sequences. If every component of the qualitative state is measured, then there can be only a single interpretation for each segment. Usually there are several, so we must generate the set of qualitative states that could give rise to the measurements.

The "one look" theory of measurement interpretation cited previously describes a solution to this problem for QP theory. We now generalize it. Call the states in the total envisionment which are consistent with the measurements represented by some global segment its *p-interps*. The possible interpretations of each global segment is exactly this collection of states. As the 1983 paper illustrated, this set may be computed via dependency-directed search over the space of possible qualitative states, pruning those which are not consistent with the measurements. If instead the total envisionment has been explicitly generated, then p-interps can be computed by table-lookup (the implications of this fact are discussed below).

However p-interps are computed, any system of reasonable complexity will give rise to many of them. Therefore it is important to prune out inconsistent interpretations as quickly as possible. Any domain-specific information applicable to the one-look case, as described in the 1983 paper, could again be useful in this context. However, when we have close data we can impose "grammatical" constraints, ruling out those p-interps which cannot possibly be part of any consistent pattern of behavior.

To impose these constraints we need to refer to the possible transitions between qualitative states contained in the total envisionment. We assume that associated with each qualitative state St is a set of *afters* which are the set of states which can be reached from St via a single transition. The following assumption is needed to apply this information:

Simplest Action Assumption: The qualitative states St_1 and St_2 which describe the behavior of two global segments S_1 and S_2 which are temporally adjacent in a close sequence (i.e., Meet(Int(S_1), Int(S_2))) are temporal successors in the total envisionment, i.e. $St_2 \in$ Afters(St_1).

In essence, this is a "compatibility constraint" applied to action. For it to be true st, our sampling time, must be small enough so that all important changes are reflected in the data. The temporal adjacency between S_1 and S_2 implies that any state which serves as an explanation for S_1 must have a transition that leads to some state which explains S_2. Similarly, any state which explains S_2 must result from some state which explains S_1. These facts can be used locally, via Waltz filtering, to prune p-interps as follows:

Given global segments S_1, S_2 s.t. Meets(Int(S_1), Int(S_2)),
  For each $St_1 \in$ p-interps(S_1) and $St_2 \in$ p-interps(S_2)
    if $\neg\exists$ $St_0 \in$ p-interps(S_2) s.t. $St_0 \in$ Afters(St_1),
      then prune St_1 from p-interps(S_1)
    if $\neg\exists$ $St_0 \in$ p-interps(S_1) s.t. $St_2 \in$ Afters(St_0),
      then prune St_2 from p-interps(S_2)

These rules must be applied to each global segment in turn until no more p-interps are pruned. Suppose for some global segment S, p-interps(S) = {}. Then either (a) the data is inconsistent or (b) the simplest action assumption is violated, either because there is more than one qualitative state required to explain a particular global segment (the hidden transition problem described previously) or the sample time st is not short enough.

Suppose the p-interps for a segment include states that are temporally adjacent, that is, for some St_1 and St_2 in p-interps(S), $St_2 \in$ Afters(St_1). Since St_1 and St_2 are in the same set p-interps(S), they must be indistinguishable with respect to the measurements provided. This is exactly how the hidden transition problem arises, and in fact is the only way it can arise – otherwise, the set of p-interps would be incomplete. Thus to find hidden transitions it suffices to extend the collection of p-

---

[4] This graph was generated by a numerical simulation program; it does not represent actual measurements. The numbers were hand-translated to Ds values.

[5] This would not be true if our system model contained an infinite number of parts. We assume such models can always be avoided.

interps to include all sequences of states from the original collection which are temporally adjacent.

Two points should be made about this pruning algorithm. First, in cases where the measurements are not very constraining the number of such sequences could grow very large. In the limiting case of no relevant measurements, the set of p-interps would correspond to the set of all possible paths and connected subparts of paths through the envisionment! We suspect such cases could arise when reasoning about a very large system with several loosely-connected components while only watching a little piece of it, and hence suggest instead a scheme combining pruning with backup for those circumstances.

Second, the algorithm can easily tolerate extra states in the sets of p-interps, but will be sensitive to missing states. These properties follow directly from the fact that states are only pruned when certain other states cannot be found. This means that gaps in the initial data will show up very rapidly, without extensive global computations.

3.1. QP example, part 2

Consider again the physical situation involving liquids discussed previously. The only processes we will be concerned with are heat flow to the liquids (if any), heat flow to the beaker, and boiling. We ignore any gasses that are produced, the possibility of the beaker melting or exploding, and any heat flow to the atmosphere. While we do not assume knowledge of the actual boiling points of water or alcohol, we assume that the boiling temperature of alcohol is lower than the boiling temperature of water. Given these assumptions, Figure 3 shows the total envisionment for the possible configurations of objects.

Since our only available measurement is temperature there is a great deal of ambiguity, as indicated by the p-interp lookup table in Figure 4. Allowing the program to apply the pruning rules, we find that after four iterations a unique solution has emerged (see Figure 5).

Figure 3 – Total Envisionment for liquids problem
The states in the picture below are divided into rows based on the contents of the beaker. A thick arrow from the burner through the beaker indicates heat flow, and small bubbles indicate boiling. A thin arrow indicates a possible transition from the state at the tail to the state at the head.

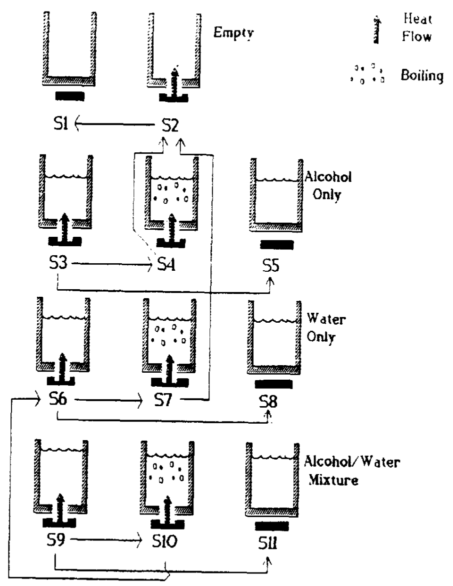

Even with very little initial data, we can conclude from this result that originally there was a mixture of water and alcohol in the beaker (S9). The mixture heated up until the alcohol started to boil (S10). After the alcohol boiled away the water heated up (S6) and began to boil (S7). After the water boiled away, the beaker heated up (S2) until thermal equilibrium was attained (S1).

4. Constructing global interpretations

Suppose the initial data is close. Then if it is correct we have a complete collection of initial hypotheses, and if the simplest action assumption is not violated and that the data is consistent, as indicated by a non-null set of p-interps for each total segment, then we have an exhaustive set of possibilities for each segment. Furthermore, the hypotheses for each segment are temporally adjacent, i.e. they are plausible candidates to follow one another in a valid description of behavior. Given these assumptions, constructing all the consistent global interpretations is simple:

Figure 4 – Table of Ds values and corresponding P-interps

Ds[T]	P-interps
0	S1, S4, S5, S7, S8, S10, S11
1	S2, S3, S6, S9

1. Select an element of the p-interps for the earliest segment.

2. Walk down the after links between p-interps, depth first.

Each such path is a consistent global interpretation.

However, close data can be hard to get. Many physically important transitions occur in an instant. For example, collisions can happen very fast; we may see a ball head into a wall and head out again without actually seeing the collision. In general we must live with sparse data. Consequently, we next describe how gaps in the data can be filled.

4.1. Filling gaps in sparse data

The procedure above can be modified to work on sparse data, although more ambiguity, and hence more interpretations, are likely.

1. Use the procedure above on all close subsequences.

2. For each gap between close subsequences, let S_1 be the segment which ends at the start of the gap, and let S_2 be the segment which starts at the end of the gap.

2.1 Select an element of p-interps(S_1).

2.2 Walk down the after links through states in the envisionment until an element of p-interps(S_2) is reached. Each such path is part of a global interpretation.

There are two cases where gaps can arise. Gaps can be small because instantaneous states have been missed, or large because the sequences are sparse. An example of a large gap is when we see a

Figure 5 – Applying Waltz filtering to P-interps

Here is the program's operation on the p-interps shown in Figure 4.

```
ATMI: Finding p-interps...
   Global Segment 1 has 4 p-interps.
   Global Segment 2 has 7 p-interps.
   Global Segment 3 has 4 p-interps.
   Global Segment 4 has 7 p-interps.
   Global Segment 5 has 4 p-interps.
   Global Segment 6 has 7 p-interps.
ATMI: Filtering p-interps...
 After 4 rounds, 27 p-interps excluded.
ATMI: Finding global interpretations...
 There is a unique global interpretation:
  (S9 S10 S6 S7 S2 S1)
The qualitative states are:
S9: water and alcohol, heat flow to beaker,
  temperature increasing.
S10: water and alcohol, heat flow to beaker,
   alcohol boiling, temperature constant.
S6: water, heat flow to beaker, temperature increasing.
S7: water, heat flow to beaker, water boiling,
  temperature constant.
S2: empty, heat flow to beaker, temperature increasing.
S1: empty, thermal equilibrium, temperature constant.
```

glass of water sitting on a table one day and come back the next day to find the glass turned over and a puddle of water on the floor. The above procedure is quite useful for small gaps, since there will be few states (usually one) between S_1 and S_2.

However, explicitly generating the set of global interpretations for large gaps can lead to combinatorial explosions. In the worse case the number of interpretations is the set of all paths through the envisionment. If the envisionment has cycles, corresponding to oscillations in behavior, the number of paths can be infinite. An alternate strategy is to use the envisionment as a "scratchpad", using the measurements to directly rule certain states in or out, and using algorithms akin to garbage collection to determine the indirect consequences of these constraints. Algorithms to do this have been implemented (see [Forbus, 1980]), and have been successfully used with the measurement interpretation program (see [Forbus, 1986]).

5. Discussion

This paper has presented a theory of interpreting measurements taken across time, illustrating its utility by extended example. The theory solves a central problem in qualitative physics, and has many potential applications. For example, this theory is useful for diagnosis problems because it provides a general ability to test fault hypotheses to see if they actually explain the observed behavior. Currently we are coding routines to automatically perform the signal/symbol transformation for several domains and generalizing the implementation to handle sparse data.

Importantly, the theory relies on very few assumptions. The small number of assumptions makes the theory applicable to many different representations and domains. The assumptions of a total envisionment and of algorithms which can provide some qualitative description of numerical parameters are very mild restrictions which most systems of qualitative physics can easily satisfy. There is no apparent reason why this theory cannot be used with device-centered models, such as [de Kleer & Brown, 1984], [Williams, 1984], or discrete-process models, such as [Simmons, 1983], [Weld, 1984], or even equation-centered models, such as [Kuipers, 1984]. In fact, we expect that the constraints on partitioning numerical data will be the same for any system that uses continuous quantities.

An interesting opportunity arises when the particular physical system is known in advance, as is typically the case when dealing with engineered systems. Current qualitative reasoning programs are often slow[6], especially when generating the entire space of possible behaviors while taking different fault modes into account. However, given a description of the structure of the system and an adequate qualitative physics, the total envisionment (or several total envisionments, representing typical fault modes) can be precomputed and preprocessed to provide a set of state tables, indexed by possible values of measurements or sets of measurements. These lookup tables, while possibly quite large, could make the interpretation process very fast. It does not seem unlikely that, given fast signal-processing hardware to perform the initial signal to symbol translation, special-purpose measurement interpretation programs which operate in real time on affordable computers might be written. As qualitative physics progresses, leading to standardized domain models and fault models, diagnostic expert systems could be automatically compiled from the structural description of a system.

5.1. Acknowledgements

Discussions with Johan de Kleer, Dan Weld, Brian Falkenhainer, Dave Waltz, and Dave Chapman led to significant improvements in both form and content. John Hogge provided invaluable programming assistance. This research is supported by the Office of Naval Research, contract No. N00014-85-K-0225.

6. Bibliography

Allen, J. "Maintaining knowledge about temporal intervals", TR-86, Department of Computer Science, January 1981

Bobrow, D., Ed. *Qualitative Reasoning about Physical Systems.*, MIT Press, 1985.

de Kleer, J. "Qualitative and quantitative knowledge in classical mechanics", MIT Artificial Intelligence lab TR352, December, 1975

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

Forbus, K. "Spatial and qualitative aspects of reasoning about motion" Proceedings of the first annual conference of the American Association for Artificial Intelligence, August 1980.

Forbus, K. "Qualitative reasoning about physical processes" Proceedings of the seventh International Joint Conference on Artificial Intelligence, August, 1981

Forbus, K. "A study of qualitative and geometric knowledge in reasoning about motion" MIT Artificial Intelligence lab TR-615, February, 1981

Forbus, K. "Measurement interpretation in qualitative process theory" in Proceedings of IJCAI-8, Karlsruhe, Germany, 1983

Forbus, K. "Qualitative process theory" *Artificial Intelligence*, **24**, 1984, pp 85-168.

Forbus, K. "Interpreting observations of physical systems", Report No. UIUCDCS-R-86-1248, Department of Computer Science, University of Illinois, Urbana, Illinois.

Kuipers, B. "Commonsense reasoning about causality: Deriving behavior from structure", *Artificial Intelligence*, **24**, 1984, pp 169-204

Reddy, D.R., L.D. Erman, R.D. Fennel, and R.B. Newly "The HEARSAY speech understanding system: an example of the recognition process" IJCAI-3, 1973, pp 185-193

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

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

Weld, D. "Switching between discrete and continuous process models tc predict genetic activity" MIT Artificial Intelligence Lab TR-793, October 1984

[6] Although efficiency can be improved: Our new implementation of QP theory (QPE) is currently running 95 times faster than our old implementation (GIZMO).