

Ken Forbus
 The Artificial Intelligence Laboratory
 Massachusetts Institute of Technology
 545 Technology Square
 Cambridge, Mass. 02139 USA

Abstract

Interpreting measurements of physical systems consists in part of constructing an account of "what's happening" in terms of our commonsense physical theories. Since most systems involve change, qualitative dynamics plays a central role in such deductions. This paper presents a theory of measurement interpretation at an instant, based on Qualitative Process Theory. Appropriate notions of measurement and interpretation are defined and the computational issues involved in constructing interpretations are examined. After describing an algorithm and illustrating its use by example, possible extensions to interpreting measurements over time will be discussed.

1. Introduction

To understand what is happening in a physical system, we must explain what we observe in terms of our physical theories. Consider for example the situation illustrated in figure 1. Given that the levels in B and C are decreasing, we would conclude that there was a flow from C to B and from B to A, and that if we could measure the level in A we would find it increasing because of water flowing into it. Measurement interpretation is an important problem in Naive Physics [1,2], and solving it will be necessary if we are to build programs that help us explain, operate and repair the complex physical systems that comprise much of our technology.

Our notion of "what is happening" in a system is intimately connected with what can be called *qualitative dynamics*. Qualitative Process theory [3,4] concerns the form of qualitative dynamics theories, postulating *physical processes* as the ultimate cause of changes. Thus QP theory provides a representational framework on which a domain-independent theory of measurement interpretation can be

erected. The general problem of measurement interpretation can be split into two cases: figuring out what is happening in a system at a particular time (taking "one look") and describing what is happening over a span of time.¹ This paper presents a theory for the first case. After a brief sketch of QP theory, we define notions of measurement and interpretation, including an account of error due to limited resolution. The design considerations for algorithms that construct interpretations are examined, along with a sample algorithm. After working through an example, the issues involved in extending the theory to the interpretation of measurements over time are discussed.

2. QP Theory in Brief

We will consider a physical situation as composed of objects and the relationships between them. The continuous parameters of an object, such as temperature and pressure, are represented by *quantities*. A quantity consists of two parts, an *amount* and a *derivative*, each of which are *numbers*. The functions *A* and *D* map from quantities to amounts and derivatives respectively. Every number has parts *sign* and *magnitude*. The functions *s* and *m* map from numbers to signs and magnitudes respectively. The parts of quantities are denoted by the selectors $A_s, A_m, D_s,$ and D_m which map from a quantity to the sign of its amount, the magnitude of its amount, and so forth.

Numbers, magnitudes, and signs take on *values* at particular times; when we wish to refer to the value of a quantity or some part of it *q* at a particular time (either instant or interval) *t*, we write:


$$(M Q t)$$

which can be taken to mean "what God would measure for *q* at *t*". Signs can take on the values -1, 0, 1. For defining comparison and combination over numbers and magnitudes we will take values to be elements of \mathbb{R} , although in QP theory we will never know the actual numerical values.

In QP theory, the value of a number² is defined in terms of its *Quantity Space* - a collection of inequalities which hold between it and other numbers. Elements in a particular Quantity Space come from the conditionalized descriptions that involve that parameter. In general the value will be incomplete, which is reflected by the Quantity Space being a partial instead of a total ordering. The Quantity Space is a useful qualitative representation because processes typically start and

Fig. 1. An Example of Measurement Interpretation

All three containers are partially filled with water, and we see that the levels in B and C are decreasing. Why?



1. Simmons [5] explores the related problem of reconstructing a sequence of events which could lead to a static final state.
2. By convention, when we speak of the value of a quantity we are referring to the value of its amount.

stop when inequalities between parameters change.¹

A changing quantity is said to be *influenced*. There are two kinds of influence which can occur. A quantity can be *directly influenced* by a process or processes, in which case we will write

$$I+(Q, n) \text{ OR } I-(Q, n)$$

according to whether n , some number defined within the process, is a positive or negative influence on the quantity Q . If a quantity is directly influenced, then its derivative is the sum of the direct influences. A quantity can also be *indirectly influenced* by its value being a function of other quantities which are themselves influenced. We will specify functional dependencies between quantities by

$$R1 \propto_Q R2$$

(read " $R1$ is *qualitatively proportional* to $R2$), meaning there exists a function which determines $R1$ and is increasing monotonic in its dependence on $R2$. Note that the function implicitly specified by \propto_Q may or may not depend upon other quantities or properties. The description of a contained liquid in figure 2, for example, uses \propto_Q to describe the relationships between the parameters of a liquid in a container. \propto_{Q-} signifies the same, but with the implicit function being

Fig. 2. Describing Liquid in a Container

In Hayes' ontology of liquids the individuating criteria for a piece of liquid is being inside some piece of space. When we have some substance s in liquid form within a container c , we have an individual that is a contained liquid, as specified below.

```
individual-view contained-liquid
Individuals:
  c a container
  s a substance
Preconditions:
  ContainsSubstance(c, s)
QuantityConditions:
  A[Amount-of-in(s, c)] > ZERO
  A[Temperature(s, c)] > TMelt(s, c)
  A[Temperature(s, c)] < TBoil(s, c)
Relations:
  There is g, a piece-of-stuff
  HasQuantity(g, Amount-of)
  Amount-of(g) = Amount-of-in(s, c)
  HasQuantity(g, Level)
  Level(g)  $\propto_Q$  Amount-of(g)
  HasQuantity(g, Pressure)
  Pressure(g)  $\propto_Q$  Level(g)
;the notation will be explained shortly
```

1. This is an application of the *relevance principle* of qualitative reasoning. Qualitative reasoning about a continuous thing requires quantization of some sort to induce a finite vocabulary of symbols. The choice of quantization must be chosen to draw the distinctions required by the kind of reasoning being performed. Ignoring this principle leads to ad hoc, inadequate, and unextendable representations.

decreasing monotonic in its dependence on $R2$.

Influences provide a means of partially specifying direct and indirect effects. Figuring out the sign of a derivative for a quantity from its influences will be called *resolving* its influences. If a quantity has no influences then its ∂_s value is 0. If all influences are of the same sign then the ∂_s value is just that sign. Inequality information can provide a resolution when direct influences have conflicting signs. Conflicting indirect influences cannot be resolved within QP theory due to the abstract nature of \propto_Q . Domain- and situation- specific information is used instead. This information can take many forms, including quantitative theories and "rules of thumb", depending on what the reasoner knows and the desired precision of answers.

A process is a thing that acts through time to cause changes. A process is specified by five parts:

Individuals: Descriptions of the entities which participate in the process,
Quantity Conditions: Inequality statements and status assignments to other conditionalized descriptions which must be true for the process to be active

Preconditions: Statements other than Quantity Conditions that must be true for the process to be active. Often Preconditions will not be deducible solely within QP theory.

Relations: The relationships between the individuals which hold when the process is active.

Influences: Descriptions of what quantities of the individuals are directly influenced by the process.

Figure 3 provides an example. A collection of objects that matches the individual specifications of a process gives rise to a *Process Instance* (PI) that represents a potential occurrence of that process. A Process Instance has a *status*, which is either *Active* or *Inactive*. A PI is active whenever both the Preconditions and Quantity Conditions are true.

Preconditions are distinct from Quantity Conditions because some factors are external to the dynamics of a domain, e.g. a purely physical theory cannot predict whether or not someone will walk by and turn on a stove, although it can predict that a result of this action

Fig. 3. Process Description of Fluid Flow

This process describes one case of fluid flow (see [2] for a partial taxonomy of cases).

```
process fluid-flow
Individuals:
  s a contained-liquid
  d a contained-liquid
  path a fluid-path, Fluid-Connection(s, d, path)
Preconditions:
  aligned(path)
QuantityConditions:
  A[Pressure(s)] > A[Pressure(d)]
Relations:
  Let flow-rate be a quantity.
  flow-rate  $\propto_Q$  (A[Pressure(s)] - A[Pressure(d)])
Influences:
  I+(Amount-of(d), A[flow-rate])
  I-(Amount-of(s), A[flow-rate])
;A fluid path is aligned only if either
;it has no valves or every valve is open
```

will be that the water sitting in a kettle on top of it will soon boil.

The statements in the Relations and Influences field hold whenever the process is active. The Relations field contains information about functional dependencies induced by the process as well as any new entities which it introduces by virtue of being active. The Influence field specifies direct influences, as discussed above.

A *Process Vocabulary* consists of the processes that form a domain's dynamics. The collection of active PI's at a particular time in a situation is called its *Process Structure*. A principle tenet of QP theory is that only processes directly influence quantities and that functional dependencies induced by processes (and other conditionalized descriptions, see below) are the causes of indirect changes. By making closed world assumptions over a Process Vocabulary and situation description, we are thus justified in reasoning by exclusion.¹

Individual Views represent states of objects and objects whose existence depends on the values of quantities, such as the Contained-Liquid description introduced above. Individual Views are specified by their Individuals, Preconditions, Quantity Conditions, and Relations just as processes are. Similarly, the *View Vocabulary* describes the Individual Views of a domain, and the *View Structure* describes the collection of view instances actually true in a situation at a particular time.

3. Measurement Interpretation

A theory of measurement interpretation must specify what can be measured, what an interpretation of those measurements is, and how to compute them from measurements. We will examine each in turn.

3.1 Measurements

First we must specify what kinds of things can be observed in principle and then add further conditions to specify what can be observed in fact. We will start by assuming that a collection of individuals is known. The closed world assumption that these individuals are the only (relevant) individuals will be called the *Armchair Assumption*. We will also assume the existence of a partial decision procedure for determining whether or not relationships defined outside QP theory (such as *FluidConnection*) hold, in order to confirm Preconditions. To state that we have ascertained whether or not a fact is true via observation, we will write:

$$\text{Observed}(\langle \text{fact} \rangle, \langle \text{time} \rangle, M)$$

where M is the instrument (such as our eyes) used in the observation.

Within the QP ontology, the kinds of facts that can be observed are occurrences of processes, inequalities and the values of signs. The criteria for a type of process being observable reduces to the observability of a particular kind of quantity and the uniqueness of that process (with respect to the reasoner's Process Vocabulary) in influencing it. A change in position, for example, is by definition the result of motion. Thus whenever we see a change of position we are seeing the result of a motion process.

We will say

1. Hayes [1,2] provides convincing arguments and examples of the role of reasoning by exclusion in Naive Physics.

$$\text{Observable}(Q, B, M)$$

when quantity (or part of a quantity such as its Q_1) Q of object B can (under some conditions) be observed with instrument M . For a process instance we will say

$$\text{Observable-PI}(PI, M)$$

if PI can be observed to be active or inactive with instrument M . We will say

$$\text{Measured}(Q, B, t, M, \text{value})$$

when measuring $Q(B)$ with M at (during) time t yields the value given. In keeping with the Quantity Space representation, measuring a number or magnitude yields an inequality and measuring a sign yields one of -1, 0, or 1. Measuring derivatives will be discussed shortly.

We wish to consider a wide variety of instruments as measuring means, such as eyes and gauges. A comprehensive theory of error lies outside QP theory, but we can model the potential for error due to limited resolution. The essence of the limitation is that when two things are "very close" a particular M might not distinguish them (for signs, we are seeing if the number is "very close" to the special value ZERO). For each measuring means, object, and quantity type, let there be a function O_{min} such that two values are considered *distinguishable* if and only if the magnitude of their actual difference is greater than O_{min} . In other words,

$$\begin{aligned} &\text{Distinguishable}(Q, B, Q_1, t, M) \\ &\leftrightarrow m[(M Q(B) t) - (M Q_1 t)] > O_{min}(Q, B, M) \end{aligned}$$

O_{min} will be chosen according to the particular physics and instrument being modelled; this particular form is chosen for simplicity.¹ A measured equality might be wrong due to limited resolution:

$$\begin{aligned} &\text{Measured}(Q, B, t, M, (= Q(B) Q_1)) \\ &\Rightarrow [(M Q(B) t) = (M Q_1 t)] \\ &\vee \neg \text{Distinguishable}(Q, B, Q_1, t, M) \end{aligned}$$

and if we measure a difference, then there really is a difference:

$$\begin{aligned} &\text{Measured}(Q, B, t, M, (> Q Q_1)) \Rightarrow \\ &[(M Q(B) t) > (M Q_1 t)] \\ &\wedge \text{Distinguishable}(Q, B, Q_1, t, M) \end{aligned}$$

A similar statement can be made for $<$. In measuring signs we are examining inequalities with $Q_1 = \text{ZERO}$, so lack of resolution will show up as a sign value of 0.

Measuring change is particularly important. First consider changes in a quantity over an interval. We must distinguish the values of the same quantity measured at two different times, so the relation we are looking for depends on the quantity type, the object, the measuring means, and an interval. We will say

$$D\text{-distinguished}(Q, B, I, M)$$

(read "differentially distinguished") exactly when

1. In particular, this form of O_{min} is simpler than *just noticeable difference* in psychophysics, because the latter also depends on the value of the quantity. A taxonomy of possible forms for O_{min} is outside the scope of this paper.

$$m[(M Q(B) \text{ start}(I)) - (M Q(B) \text{ end}(I))] > \text{Omin}(Q, B, M)$$

If we adopt Allen's ontology for time [6], then an instant is simply a very short interval. Thus our criteria for observing changes over intervals can serve for measuring derivatives. However, capturing lack of resolution becomes more complicated:

$$\begin{aligned} \text{Measured}(Ds[Q], B, I, M, 0) \\ \Rightarrow (\forall t \in \text{during}(I) (M Ds[Q(B)] t) = 0) \\ \vee \neg D\text{-distinguished}(Q, B, I, M) \\ \vee \neg \text{Constant-Sign}(D[Q(B)], I) \end{aligned}$$

Even when the measured value is non-zero, the sign may not have been constant over the interval:

$$\begin{aligned} \text{Measured}(Ds[Q], B, M, I, \langle 1 \text{ or } -1 \rangle) \\ \Rightarrow [(\forall t \in \text{during}(I) (M Ds[Q(B)] t) = \langle 1 \text{ or } -1 \rangle) \\ \vee \neg \text{Constant-Sign}(D[Q(B)], I)] \\ \wedge D\text{-distinguishable}(Q, B, I, M) \end{aligned}$$

This extension allows us to say, for example, that while we cannot immediately see the effect of evaporation on the level of water in a glass, if we looked longer we could.

3.2 Interpretations

An interpretation must explain what is causing the changes that are occurring (including the special case of nothing changing). In QP theory processes are the only causes of changes, so an interpretation will include assumptions about the status of the process instances (PI's) that occur between the individuals. Since more than one process can influence a quantity, interpretations must also include assumptions concerning influence resolutions. An interpretation must be *internally consistent*, *externally consistent*, and *sufficient*. Internally consistent means an interpretation assigns at most one status to any PI and at most one d_s value to any quantity. Externally consistent means that the status assignments and d_s values assigned are consistent with the measurements. Sufficient means that every measured d_s value is explained.¹

Some additional structure on interpretations will prove useful. A *Unit Cause Hypothesis* (UCH) is a partial interpretation that forces the assignment of a d_s value consistent with a measurement. Any interpretation which satisfies the three criteria above will be a collection of UCHs, one for each d_s measurement, that is internally consistent. The *P-influencers* of a quantity is the set of process instances that can possibly influence that quantity, directly or indirectly. The *Influencers* of a UCH is the subset of the P-influencers that are active in that UCH. In addition to the status assumptions that determine the influencers, a UCH with conflicting influences must include an assumption about their resolution. As noted above, for direct influences this will take the form of an inequality between (perhaps sums of) the influences. To state the resolution of conflicting indirect influences we will say

$$\text{Resolved}(\langle \text{quantity} \rangle \langle \text{influencers} \rangle \langle \text{value} \rangle)$$

3.3 Computational Issues

There are two possible ways to organize the search for interpretations. One way is to search through the possible UCH's for each measurement to find a globally consistent collection. Finding the possible UCH's for a quantity is simple. The set of p-influencers can be computed from the process descriptions associated with the PI's, and each possible subset of influencers can be checked to see if it can be resolved consistently with the measured value. However, the number of UCH's can be quite large. Suppose we measure a quantity and find it is increasing. Then if we have p PI's that can provide a positive influence and n that can provide a negative influence, there are

$$(2^p - 1) \cdot (2^n)$$

possible UCHs. In practice this number will be much smaller, since the PI's are usually not independent. For example, a fluid path cannot have flows going in both directions at once because their Quantity Conditions would conflict (see figure 3). The number of consistent interpretations will almost always be much smaller than the product of the number of UCHs since processes typically influence more than one quantity, providing mutual constraint. These facts suggest that we organize the search around the space of status assignments to PIs instead. If we wish a *total* interpretation we can use the entire collection of PIs, but if we want a minimal interpretation to explain the measurements we can just use the union of the P-influencers for the measured d_s values. Any collection of status assignments that cannot be consistently extended by assumptions about influence resolutions to provide a UCH for each d_s measurement can be thrown out, and each extension found is a valid interpretation.

Several kinds of knowledge can be used to prune the search space. PIs that correspond to observable processes could have their status determined directly by observation, or indirectly by ascertaining the truth of their Preconditions and Quantity Conditions. For example, if we can see that a valve in a fluid path is closed, then that fluid path is not aligned and no flows can occur through it.

Once a collection of status assumptions is chosen, it must be extended to form a collection of UCHs. There are several ways to accept or rule out a UCH. If the set of influences can be resolved then the UCH will stand or fall according to whether or not the resolved d_s value and the measured value agree. Again, this can require domain-specific information; we do not expect that evaporation will immediately cancel out the effect of pouring water into a cup. Distinguishability provides a means of ruling out small changes. For example, we can say:

$$\begin{aligned} \forall w \in \text{contained-liquid} \forall pi \in \text{process-instance} \\ \text{Influencers}(\text{Level}(w)) = pi \wedge \text{Process}(pi) = \text{Evaporation} \\ \Rightarrow \neg D\text{-Distinguished}(\text{Level}, w, \text{eyeball-time}, \text{eyes}) \end{aligned}$$

which will rule out evaporation as the sole explanation for why we are

1. By contrast, an interpretation in deKleer's QUAL is a collection of device states and incremental changes in quantities, the latter assumed to occur sequentially in "mythical time" [7]. Despite profound ontological differences, the principles defining interpretations presented here are inspired by his work

seeing the level of water in a glass fall.¹

What if no consistent interpretation exists? Following [8], we view the analysis as relying on simplifying assumptions that must be re-analyzed in such cases. The assumptions, ordered in increasing certainty, are:

1. Any facts used in pruning are correct
2. The measurements are correct
3. The Armchair Assumption is correct
4. The Process Vocabulary is complete and correct

Other orderings are of course possible.²

3.4 An Algorithm

The prescriptions above are combined into the algorithm in figure 4. MI-1 constructs all interpretations for a set of measurements, given a collection of individuals, relationships, and a Process Vocabulary. MI-1 can easily be modified to produce just one interpretation by stopping the search after the first interpretation has been generated. To evaluate whether or not a proposed interpretation is consistent (as would be necessary for training, see [9]) the proposed interpretation can be included in the list of OFACTS. Of course, there are other algorithms that are consistent with the theory. For example, the mix of search to information gathering will depend on the particular domain and available instruments. A version of MI-1 has been implemented and successfully run.

4. Example

Let us return to the initial example to illustrate the ideas. Figure 5 provides the initial description of the situation. We assume that contained liquids (see figure 2) exist in the containers A, B, and C, called WA, WB, and WC respectively. The Process Vocabulary will consist of fluid flow (see figure 3), and P1 and P2 are assumed to be fluid paths that form fluid connections between WA and WB, WB and WC respectively. We want to find causes for the drop in the levels of WB and WC.

Following MI-1, there are 4 process instances of fluid flow, corresponding to flow in each direction of each fluid path. ML consists only of D_s values, so OBS = ML and OFACTS is empty. Figure 6 summarizes the results so far. Assume we have no extra information about the domain or the situation. Since each set of P-influencers

1. Informal observations indicate that people appear to use the following *Ineffectuality heuristic* - if a PI's result is not distinguishable, assume it isn't acting. The intuition appears to be that its effect won't make that much of a difference anyway (unless the physical structure of the situation leads you to believe there are a lot of them!). This heuristic prunes the search space of PIs enormously, and it seems likely that correct use of this assumption is a mark of an expert in a domain.

2. Ultimately a global order on categories of facts will be inadequate, since our strength of belief seems to vary on items within a category. For example, when the measuring means is indirect and the domain familiar, we often trust our theories more than the measurements. The opposite is true if the measurements are direct (sensory) and the domain unfamiliar.

Fig. 4. Measurement Interpretation Algorithm

```

;Let ML be the list of measurements.
;IS be the set of individuals which comprise
;   the situation.
;IV be the vocabulary of individual views
;PV be the Process Vocabulary

Procedure MI-1

1. Compute all instances of Individual Views (IVI's) and
   Process Instances (PI's) for the individuals in IS by finding
   matches with elements of IV and PV.

2. Partition ML into two sets, OBS for assertions of  $D_s$  values
   and OFACTS for all others. For each quantity in OBS compute
   the set of p-influencers.

3. Make status assignments to PI's (and IVI's) wherever
   possible, using OFACTS and domain specific knowledge.
   For any quantity in OBS whose p-influencers set
   contains a single element.

   - If the sign of the influence is consistent with the
     observed  $D_s$  value, the PI is ACTIVE.

   - If the sign of the influence is inconsistent with the
     observed  $D_s$  value, then there is a global inconsistency.

4. Perform a dependency directed search over the set of status
   assignments remaining, using the following criteria to
   determine inconsistency:

   - If inconsistent assumptions about Preconditions and
     Quantity Conditions result, or

   - If no consistent UCH can be constructed for some quantity
     in OBS (either status assumptions force an incorrect  $D_s$ 
     assignment or either situation-specific or domain-specific
     knowledge rule out the correct assignment or the assumptions
     of another UCH rule it out), then the status assignments are
     inconsistent.

   - If not ruled out, the collection of status assumptions
     along with the assumptions made in constructing the UCH's
     comprise a valid interpretation.

```

Fig. 5. Initial Facts

```

At time  $t_0$ 
  Contained-Liquid(WA) ;the liquid in container A
  Contained-Liquid(WB)
  Contained-Liquid(WC)
  Fluid-Path(P1)
  Fluid-Path(P2)
  FluidConnection(WA, WB, P1); assume no
  FluidConnection(WB, WA, P1); check valves
  FluidConnection(WB, WC, P2)
  FluidConnection(WC, WB, P2)

ML = (M  $D_s$ [Level(WB)]  $t_0$ ) = -1.
      (M  $D_s$ [Level(WC)]  $t_0$ ) = -1

```

contains more than one PI, step 3 yields no results. We now must search over the status assignments for FF1-4. FF3 and FF4 cannot both be active because they presuppose different orderings between $A[\text{Pressure}(WB)]$ and $A[\text{Pressure}(WC)]$. If FF3 alone is active then $D_s[\text{Level}(WC)]$ would be 1, which contradicts the measured value. FF4 being active results in an influence resolution consistent with the

Fig. 6. Intermediate Descriptions of the Situation

Process Instances:

```

FF1: Fluid-Flow(WA, WB, P1)
FF2: Fluid-Flow(WB, WA, P1)
FF3: Fluid-Flow(WB, WC, P2)
FF4: Fluid-Flow(WC, WB, P2)

```

```

OBS = (M Ds[Level(WB)] t0) = -1,
      (M Ds[Level(WC)] t0) = -1

```

```

P-Influencers(Level(WB)) = FF1, FF2, FF3, FF4
P-Influencers(Level(WC)) = FF3, FF4

```

measured value, providing the only consistent UCH for the measurement of $D_s[\text{Level}(\text{WC})]$. Given that FF4 must be active, FF1 cannot be, since it would result in $D_s[\text{Level}(\text{WB})]$ being 1, contradicting the measured value of -1. FF2 must be active, because if it wasn't the sole influence on $D_s[\text{Level}(\text{WB})]$ would be FF4, again contradicting the measured value. In addition, the flow rate of FF2 must be greater than the flow rate of FF4 to explain the measurement. So the only consistent interpretation is:

```

Status(FF1, Inactive)
Status(FF2, Active)
Status(FF3, Inactive)
Status(FF4, Active)
flow-rate(FF2) > flow-rate(FF4)

```

We would also predict that $D_s[\text{Level}(\text{WA})] = 1$, since FF2 is the only influence on Amount-of(WA). Suppose however that we measure $D_s[\text{Level}(\text{WA})]$ to be -1. There is no consistent interpretation of these measurements, so one of the assumptions underlying the analysis is wrong and must be retracted. Pragmatically, the Armchair Assumption is the best candidate for retraction since tanks do leak - but that decision lies outside the range of QP theory.

5. Discussion

A natural extension of this theory is interpreting measurements over time. An interpretation would be generalized to a *history* [2], and measurements would correspond to partial information about this history. The description of measurements remains unchanged, the only difference being that differential distinguishability will be used over significant intervals. There is an additional problem of *segmentation*, finding intervals where the Process Structure does not change. A heuristic is to use changes in D_s values as the boundaries, since these must correspond to changes in the resolving of influences. Additional divisions may be necessary, because changes in unobserved quantities may take time to propagate to distinguishable changes in observed quantities - for example, a stove may be on for some time before you deduce that fact by seeing steam pour out of a kettle on top of it. The theory described here could be used to build interpretations for what is occurring during each episode implied by the boundaries. Because the episodes are connected, the interpretation for any particular episode has to be consistent with the interpretations for the ones around it. The pruning constraints and heuristics described above still hold - even if we watch for five minutes, evaporation still won't empty a drinking glass. QP theory also imposes an additional constraint

on the connected episodes: the interpretation for each episode has to correspond to a Process Structure implied by some Limit Hypothesis (see [4]) for the Process Structure implied by the interpretation of the episode before it.

Interpreting observations in terms of physical theories is an important problem in Naive Physics, and the theory presented here is a first step towards solving that general problem. Using Qualitative Process Theory as the representational framework provides a fairly natural notion of what an interpretation is (what processes are acting, with what net effect) and ensures that the theory and any algorithms based on it can be used in more than one domain. While many domain-specific facts will be needed in any practical system that performs measurement interpretation, QP theory provides a common language in which to express at least part of the information, and provides a constrained role for other kinds of information (resolving conflicting influences, determining Preconditions, Quantity Conditions, and distinguishability).

6. Bibliography

- [1] Hayes, Patrick J. "The Naive Physics Manifesto" in *Expert Systems in the Microelectronic age*, edited by D. Michie, Edinburgh University Press, May 1979
- [2] Hayes, Patrick J. "Naive Physics 1 - Ontology for Liquids" Memo, Centre pour les etudes Semantiques et Cognitives, Geneva, 1979
- [3] Forbus, K. "Qualitative Reasoning about Physical Processes" Proceedings of IJCAI-7, 1981
- [4] Forbus, K. "Qualitative Process Theory" MIT AI Lab Memo No. 664, February 1982, Revised June, 1983
- [5] Simmons, R. "Spatial and Temporal Reasoning in Geologic Map Interpretation" Proceedings of the National Conference on Artificial Intelligence, 1982
- [6] Allen, James "Maintaining Knowledge about Temporal Intervals" TR86, Computer Science Department, University of Rochester, January 1981
- [7] de Kleer, J. "Causal and Telcological Reasoning in Circuit Recognition" TR-529, MIT AI Lab, Cambridge, Massachusetts, September 1979
- [8] Davis, R., Shrobe, H., Hamscher, W., Wieckert, K., Shirley, M., and Polit, S. "Diagnosis Based On Description of Structure and Function" Proceedings of the National Conference on Artificial Intelligence, 1982
- [9] Brown, J. S., Burton, R. R., and Bell, A. G., "SOPHIE: A Sophisticated Instructional Environment for Teaching Electronic Troubleshooting (An example of AI in CAI)" Final Report, BBN AI Report No. 12, BBN, Cambridge, Mass, 1974