

Interpreting Observations of Physical Systems

KENNETH D. FORBUS

Abstract—An unsolved problem in creating diagnostic expert systems is generating a qualitative understanding of how the system is behaving from raw data, especially numerical data taken across time. Yet automating this critical step is necessary for building the next generation of expert systems. The theory described provides a means of interpreting observations made of a physical system across time in terms of qualitative theories. Importantly, the theory is *ontology-independent* as well as *domain-independent* in that it only requires 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 is illustrated step by step with two extended examples, one involving qualitative process theory and the other involving a qualitative state vector representation of motion. The performance of an implementation of the theory is also illustrated.

I. 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, especially, 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.

Currently, a great deal of interest exists in applying qualitative physics to engineering tasks such as diagnosis (e.g., the articles in [2]). 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, called *ATMI*. The *ATMI* 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 (perhaps on demand);

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

Consequently, the theory should apply to any system of qualitative physics. The theory is analogous to artificial intelligence (AI) models of speech understanding (e.g., [13]). 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 the theory presented here, 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.

Overview

The foregoing framework suggests a collection of problems that must be solved by any account of measurement interpretation.

- 1) How should the initial data be segmented?
- 2) How are interpretations constructed for each segment?
- 3) How can global constraints be applied to prune the local interpretations?
- 4) How should global interpretations be constructed?
- 5) How should gaps in the initial data be handled?
- 6) How should recovery from inconsistent data be effected?

The following theory answers questions 2–5. Our goal 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 answer question 1. We demonstrate how the theory can be instantiated using extended examples from two different domains, represented using two distinct ontologies. The first example concerns motion through two-dimensional space and is represented using the *qualitative state vector ontology* [5], [6], [8]. The second example concerns state changes in fluids and is represented using *qualitative process theory* [7], [10]. Some familiarity with these representations would be useful but is not essential. The *ATMI* theory has been implemented, and the performance of the implementation on these examples is demonstrated.

Manuscript received December 15, 1985; revised September 30, 1986 and January 25, 1987. This work was supported in part by the Office of Naval Research, Contract N00014-85-K-0225.

The author is with the Qualitative Reasoning Group, Department of Computer Science, University of Illinois, 1304 W. Springfield Avenue, Urbana, IL 61801.

IEEE Log Number 8714379.

The next section provides a vocabulary for describing the initial data and places constraints on the segmentation process. Section III generalizes the earlier theory of interpreting measurements taken at an instant [9] based on qualitative process (QP) theory and shows how the envisionment can be used to locally prune interpretations of segments.¹ Section IV illustrates how global interpretations are constructed and how even incomplete information can provide useful constraints on possible future behaviors. Finally, we discuss the implications of the theory, including how it might be extended to handle noisy data and how we plan to extend the implementation.

II. INPUT DATA AND SEGMENTATION

First we describe the kinds of inputs the theory handles. For simplicity we will assume a function *time*, which maps measurements to real numbers, and assume that the duration of an interval is simply the difference between the times for its start and end points. We also assume that the temporal relationships described in [1] may be applied to intervals. These relationships include:

Meet(I_1, I_2) means " I_2 starts directly after I_1 , with no time in between,"

Starts(I_1, I_2) means " I_1 and I_2 start at the same time,"

Finishes(I_1, I_2) means " I_1 and I_2 end at the same time."

We will say

$$\text{Observable}(\langle p \rangle, \langle i \rangle)$$

when property $\langle p \rangle$ can be observed in principle by instrument $\langle i \rangle$. To say that we can measure the level of water in a can with our eyes, we would write²

$$\text{Observable}(A[\text{Level}(C-S(\text{water, liquid, can}))], \text{eyes}).$$

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. We will say

$$\text{Measured}(\langle p \rangle, \langle v \rangle, \langle t \rangle, \langle i \rangle)$$

to mean that property $\langle p \rangle$ takes on value $\langle v \rangle$ at time $\langle t \rangle$, as measured by instrument $\langle i \rangle$.³ To say that we measured the level in the can to be 5 cm 12 s after an experiment

¹The term "envisionment" has been used in two distinct senses. In both cases it means, roughly, "generating all the qualitatively distinct possible behaviors of a system." However, sometimes it refers to "all behaviors possible from some given initial state," and sometimes it refers to "all behaviors for each possible initial state." We will refer to the first usage as *attainable* envisionments, and the second usage as *total* envisionments. In this paper we will only be concerned with total envisionments.

²The first argument in this example 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.

³This definition of *Measured* is more general than the definition introduced in the 1983 theory [9].

started, we would write, ignoring units,

$$\text{Measured}(A[\text{Level}(C-S(\text{water, liquid, can}))], \\ 5, 12, \text{eyes}).$$

For simplicity, we will use these same conventions to define observations of arbitrary facts by interpreting them as functions whose range is the set {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 are given some "grain" on time st such that events which occur in durations shorter than st will not be considered relevant.⁴ Then two types of measurement sequences must be considered:

close: the data are complete in the sense that over the total interval of interest measurements are separated by durations no larger than st ;

sampled: temporal gaps exist in the data with durations larger than st .

Given the 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.

A. 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 in some sense. A symbolic property is constant over an interval if its value is identical for all measurements within that interval. Notice that signs of derivatives, as represented in QP theory, 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 will, of course, depend on the choice of domain representation and ontology.⁵ All that we require for this theory to be applicable 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.

⁴The problem of instantaneous events will be discussed in Section IV-A.

⁵To 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.

Once a translation to qualitative values has been provided for numerical parameters, segmentation becomes simple. However, these segments cannot necessarily be identified with a single qualitative state. First, the qualitative value may be partial, as noted earlier. 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, as will be discussed later.

Each segment formed in a measurement sequence covers a nonempty collection of data, and since the data are temporally well ordered, a maximum and minimum time will be associated with each data set. Let the minimum time be the start time and the maximum time be the end time. We define two intervals I_1 and I_2 to be *adjacent* if no interval is 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 will also say that the intervals Meet, as defined in Allen's time logic. The function *Int* will be used to map 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($\text{Int}(GS_k), \text{Int}(S_{i,j})$) 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($\text{Int}(GS_k), \text{Int}(S_{i,k})$) 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 will span the largest possible interval where qualitative values are constant. These constraints ensure that global segments are good candidates for explanation by a single qualitative state.

B. Example: The Bouncing Ball World

Our first example is the *bouncing ball world*, a domain which highlights several issues involving reasoning about motion through space. A scenario in this world consists of a specification of a set of straight surfaces and the initial position of one or more balls. Geometrically, the balls are modeled as point masses. A program called FROB was




Fig. 1. Typical scenario in bouncing ball world. Consists of some surfaces and one or more balls. FROB accepts both qualitative and quantitative information about properties of balls and their motions, checks consistency of information, and can use it to answer questions about where they can end up and whether or not they might collide.




Fig. 2. Typical place vocabulary for input scenario of Fig. 1. Notice that free space has been explicitly represented by breaking it up into several pieces, within which motion is easily described. SPACE-REGION'S (denoted R) are pieces of free space. FREE-BORDER'S (denoted F) mark borders of space regions that are within diagram. BORDER'S (denoted B) indicate edges of the diagram. FROB loses all interest in balls which leave diagram, assuming that they never return. SURFACE'S (denoted S) indicate the solid surfaces specified by user. To avoid clutter, we will label elements only as necessary in future diagrams.

implemented to reason about the possible motions of balls, their eventual fate (e.g., "could this ball be trapped in the well?"), and whether or not balls could collide [5], [6], [8]. Fig. 1 illustrates a typical scenario of the bouncing ball world.

The initial diagram is represented as a *metric diagram*, a combination of numerical and symbolic information. The metric diagram provides some of the advantages people have when reasoning geometrically with a diagram. To reason about space qualitatively, the initial diagram must be used to compute a *place vocabulary* that quantizes the free space of the diagram into "interesting" pieces, like the inside of a well. Since all the moving objects in the bouncing ball world have the same shape, the same place vocabulary is appropriate for each of them. The particular principles used to construct place vocabularies in this world are described in [6]. Here we only note that free




Fig. 3. How qualitative state vectors are graphically depicted. 1: (FLY SR3 (LEFT UP)). 2: (COLLIDE S11 (RIGHT DOWN)). 3: (STOP S13 NIL). 4: (CONTINUE S49 UP).


space (and its boundaries) is broken up into nonoverlapping pieces. Let SPACE-REGION, FREE-BORDER, BORDER, and SURFACE be predicates on places that are true if the place is a region of free space, a border between two regions of free space, part of the edge of the diagram, or a piece of a surface, respectively. Fig. 2 illustrates a typical place vocabulary.

The state of a ball can be represented quantitatively (i.e., by numerical coordinates for its position and velocity) or qualitatively. The *qualitative state vector* for motion consists of the type of activity, where it is occurring, and what direction it is occurring in. The type of activity ranges over a small set of symbols, such as FLY, COLLIDE, CONTINUE, and STOP. Location is specified by an element of the place vocabulary. The direction is specified with respect to the gravity vertical, i.e. LEFT, UP, (LEFT UP), and so forth. Fig. 3 shows a sampler of qualitative states and how they can be depicted graphically.

There are two kinds of numerical parameters that can easily be measured in the bouncing ball world: coordinates and velocities. Ignoring obstructions, these will always be observable:

$$\begin{aligned} \forall B \in \text{ball} \forall t \in \text{time} \\ & [\text{Observable-at}(X(b), \text{eyes}, t) \\ & \wedge \text{Observable-at}(Y(b), \text{eyes}, t) \\ & \wedge \text{Observable-at}(V_x(b), \text{eyes}, t) \\ & \wedge \text{Observable-at}(V_y(b), \text{eyes}, t)]. \end{aligned}$$

Now let us determine the mapping between numerical and qualitative representations, which will provide the initial symbolic descriptions. Velocities are the simplest to parse into qualitative terms; the sign of the X and Y components provide a useful qualitative description because they describe which general direction the ball is going in. We assume a function *direction* which maps vectors into symbolic directions. To quantize coordinates, assume a function *place* which, given a pair of X and Y coordinates, queries the metric diagram to find out what



Metric Diagram




Fig. 4. Partitioning numerical data in FROB's world. Translating numerical data into initial symbolic description is simple in FROB's world. Particular motion generated by FROB's quantitative module is shown, along with graphical depiction of segments and global segments computed from initial symbolic data.

element of the place vocabulary that point lies in. Given these functions⁶ which map from numerical to initial symbolic information, the algorithm segments the individual measurements and calculates a global segmentation. Fig. 4 illustrates.

C. Example: Qualitative Process Theory

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 (QP) theory formalizes this intuitive notion of physical process and provides a qualitative language for differential equations that preserves distinctions required for causal reasoning [7], [10].

QP theory provides several types of measurable properties, including the truth of predicates and relations, the status of process and view instances, and, of course, information about numbers. As mentioned before, 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 (QP aficionados will recognize this as the D_s value of a quantity, which ranges over $\{-1, 0, 1\}$).

Consider the situation shown in Fig. 5. We have a beaker that either has some liquid A, some liquid B, or a mixture of both liquids inside. Suppose we place the beaker on a stove with a thermometer placed against the bottom on the inside of the beaker. In this case we can

⁶Currently, these functions are hand simulated.




Fig. 5. QP measurement interpretation problem. Suppose beaker on stove contains only liquid A, only liquid B, or mixture of both. Given only data from thermometer inside beaker, can we tell what is in there?

always measure the temperature, i.e.,

$\forall t \in \text{time}$

Observable-at($A[\text{Temperature}(\text{inside}(\text{beaker}))]$,
thermometer, t).

If we plot the temperature with respect to time, we might get the graph shown in Fig. 6.⁷

If we do not know the numerical values for the boiling points of liquid A and liquid B, 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 the following:

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 1.300002.

DS of (T INSIDE-BEAKER) is 0 from 1.400002 to 2.100001.

DS of (T INSIDE-BEAKER) is 1 from 2.2 to 4.0999985.

DS of (T INSIDE-BEAKER) is 0 from 4.1999984 to 5.799997.

DS of (T INSIDE-BEAKER) is 1 from 5.8999968 to 8.499996.

DS of (T INSIDE-BEAKER) is 0 from 8.599997 to 11.700008.

ATMI: Finding global segments...

There are 6 global segments.

III. 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 compo-

⁷This graph was generated by a numerical simulation program; it does not represent actual measurements. The numbers were hand translated to Ds values. This graph represents the common-sense model of many people; however, it is only accurate when the liquids vary widely in volatility and boiling point (see [3], [12]). Examples include solutions of methanol and water, or ethylene glycol and water. Regrettably, the suggestion that solutions of water and alcohol (ethanol) follow this pattern in [11] is incorrect.




Fig. 6. Temperature plotted as function of time. Sample plot of data that might be generated from experimental situation shown in Fig. 5.

nents,⁸ 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, this will not be the case, and 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 shall 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 [9] illustrated, this set may be computed by a 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 will be discussed later).

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 [9], could again be useful in this context. However, when we have close data, we would like to 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

⁸This would not be true if we were trying to model a system in a manner that gave rise to an infinite number of parts. We assume such models can always be avoided.

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., $\text{Meet}(\text{Int}(S_1), \text{Int}(S_2))$) are temporal successors in the total envisionment, i.e., $St_2 \in \text{Afters}(St_1)$.

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 symbolic relaxation, to prune p-interps as follows:

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

These rules must be applied to each global segment in turn until no more p-interps are pruned. Suppose for some global segment S_i , $\text{p-interps}(S_i) = \{\}$. Then either a) the data are 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 small enough.

Suppose the p-interps for a segment are temporally adjacent, that is, for some St_1 and St_2 in $\text{p-interps}(S)$, $St_2 \in \text{Afters}(St_1)$. Since St_1 and St_2 are in the same set $\text{p-interps}(S)$, 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-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 sequences introduced to solve the hidden transition problem 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 complicated cases could easily arise when observing a small fragment of a large system with several loosely connected components. We suggest a scheme combining pruning with backup for those circumstances, generating connected p-interp sequences as required.

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.

A. Example: Bouncing Balls

Given a quantitative state, the differential equations for motion can be solved to provide predictions of future states. Similarly, qualitative simulation laws exist which, given a qualitative state vector, will generate the set of qualitative state vectors which can arise from it. These laws can be used to generate attainable envisionments by applying them recursively to the results obtained by applying them to an initial state. Generating total envisionments is a two-step procedure. First, explicitly generate the space of legal qualitative state vectors by combining the possible components, such as position and direction. Second, use the simulation laws to find out what transitions are possible between the different parts of the qualitative vector space.

Computing p-interps for the bouncing ball world is straightforward. As with any domain, the state descriptions which comprise the envisionment can be inverted to produce tables for rapidly retrieving candidate states corresponding to particular observed values. However, the bouncing ball world is particularly simple. Since the elements of the place vocabulary in the bouncing ball world do not overlap, the mapping from measured positions in the metric diagram to places will be unique. Furthermore, the type of place and direction uniquely determine the type of action that occurs. For example, if the ball is at a surface and its direction is into the surface, then a collision is occurring. If its direction is away from the surface, then the ball is flying. Thus intersecting the states suggested by these two measurements always yields a single p-interp. The state path found for the measurements presented earlier is illustrated as follows.

ATMI: Filtering p-interps...

Round 1, 0 p-interps excluded.

1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1

After 0 rounds, 0 p-interps excluded.

ATMI: Finding global interpretations...

There is a unique global interpretation:

(SEQ29 SEQ67 SEQ69 SEQ31 SEQ106 SEQ22 SEQ21
 SEQ20 SEQ117 SEQ2SEQ52 SEQ57 SEQ8 SEQ50 SEQ48
 SEQ6)

The qualitative states are:

SEQ29: Flying left and down in SREGION3.

SEQ67: Colliding with SEGMENT9, heading left and down.

SEQ69: Flying left and up from SEGMENT9.

SEQ31: Flying left and up in SREGION3.

SEQ106: Flying through SEGMENT44 going left and up.

SEQ22: Flying left and up in SREGION2.

SEQ21: Flying left in SREGION2.

SEQ20: Flying left and down in SREGION2.

SEQ117: Passing through SEGMENT41 going left and down.




Fig. 7. Total envisionment for liquids problem. States in picture are divided into rows based on what is inside beaker in them. Thick arrow from burner through beaker indicates heat flow, and small bubbles indicate boiling. Thin arrow indicates possible transition from state at tail to state at head. We ignore any gasses produced, any possible changes in nature of beaker, and any heat transfer to surroundings. We also assume that boiling temperature of liquid B is less than that of liquid A, although we will not assume any particular numerical values for either temperature.

- SEQ2: Flying left and down in SREGION0.
- SEQ52: Colliding with SEGMENT12, heading left and down.
- SEQ57: Flying right and down from SEGMENT12.
- SEQ8: Flying right and down in SREGION0.
- SEQ50: Colliding with SEGMENT11, heading right and down.
- SEQ48: Flying right and up from SEGMENT11.
- SEQ6: Flying right and up in SREGION0.

B. Example: QP Theory

Consider again the physical situation involving liquids presented in Fig. 5. The only processes we will be concerned with are heat flow to the liquids (if any), heat flow to the beaker, and boiling. We will ignore any gases 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 liquid A or liquid B, we will assume that the boiling temperature of liquid B is lower than the boiling temperature of liquid A. Given these assumptions, Fig. 7 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 Table I. The p-interps for the global segments described previously, along with the transitions linking states in adjoining segments, are shown in Fig. 8. Allowing the program to apply the pruning rules,

TABLE I
Ds VALUES AND CORRESPONDING P-INTERPS^a

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

^aThe temperature can only be increasing or constant, so all the states divide into two classes for this scenario.

Initial Hypotheses

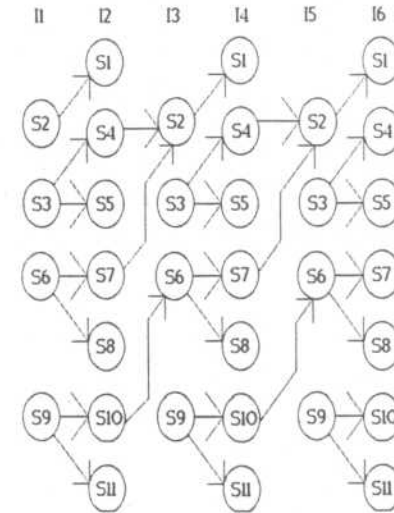


Fig. 8. Initial p-interps for QP liquids problem. Each column represents p-interps found for particular global segment from Fig. 7. Temporal ordering of segments runs from left to right. The transitions between states, indicated by arrows, represent "grammatical" constraints between them. Any unbroken path of states from left to right constitutes consistent global interpretation.

we find that after four iterations a unique solution has emerged:

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...

- Round 1, 21 p-interps excluded.
4 3 2 1 1 1
- Round 2, 2 p-interps excluded.
3 3 1 1 1 1
- Round 3, 2 p-interps excluded.
3 1 1 .1 1 1
- Round 4, 2 p-interps excluded.
1 1 1 1 1 1
- Round 5, 0 p-interps excluded.
1 1 1 1 1 1

After 4 rounds, 27 p-interps excluded.

ATMI: Finding global interpretations...

There is a unique global interpretation:
(S9S10S6S7S2S1)

The qualitative states are:

S9: liquid A and liquid B, heat flow to beaker, temperature increasing.

S10: liquid A and liquid B, heat flow to beaker, liquid B boiling, temperature constant.

S6: liquid A, heat flow to beaker, temperature increasing.

S7: liquid A, heat flow to beaker, liquid A boiling, temperature constant.

S2: empty, heat flow to beaker, temperature increasing.

S1: empty, thermal equilibrium, temperature constant.

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

IV. CONSTRUCTING GLOBAL INTERPRETATIONS

Suppose the initial data are a close sequence. Then if the data are correct, we have a complete collection of initial hypotheses, and if the simplest action assumption is not violated (i.e., we have sufficient temporal resolution so that nothing is happening "under the table") and the data are consistent as indicated by a nonnull 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 assumptions is simple.

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 come by. 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 will have to live with sparse data. Consequently, we next describe how gaps in the data can be filled.

A. Filling Gaps in Sparse Data

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

1) Use the procedure 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 $p\text{-interps}(S_1)$.

2.2) Walk down the *after* links through states in the envisionment, depth first, until an element of $p\text{-interps}(S_2)$




Fig. 9. Example involving sparse data. Total envisionment for original scenario FROB scenario. Suppose we know only that ball was originally in R3—rightmost space region—and that sometime later ball is seen in R1—the leftmost space region. Envisionment can be pruned to reflect constraints imposed by these measurements.

is reached. Each such path is part of a global interpretation.

Two cases exist 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 glass of water sitting on a table one day and coming back the next day to find the glass turned over and a puddle of water on the floor. The procedure described previously 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, which has been implemented for FROB, is described in the next section.

B. Example: Bouncing Balls

Consider the scenario in Fig. 9. Suppose all we know is that the ball was in R3 first, and later is in region R1. An infinite number of paths exist through the envisionment consistent with these measurements, corresponding to all the ways a ball might bounce around inside a well and finally bounce out again heading left. Clearly, explicit generation must be avoided in this case.

FROB uses the envisionment as a "scratch pad" to represent the consequences of assumptions about motion. Importantly, the consequences of the ATMI theory can be cast in this form. In particular, the p-interps indicate at what places the ball must appear. Recall that the envisionment represents all possible behaviors. The effect of any additional assumptions or data will be to eliminate some of these possibilities, which can be represented by pruning the envisionment. Simple algorithms, akin to those used in garbage collection (see [6] for details), determine the indi-




Fig. 10. Sparse measurements can provide useful constraints. Envisionment can be pruned to reflect constraints imposed by measurements. Comparing diagram against Fig. 12 illustrates how many possibilities can be pruned even with sparse data.




Fig. 11. Facts about object can provide still more constraint. Suppose we also find out that ball is inelastic. Then this information, combined with previous sparse measurements, allows FROB to infer that ball must not have gone inside well, and never will.

rect consequences of these assumptions. In this case, any state which cannot be part of a motion that includes the places which the ball must go through can be pruned.

Returning to our example, notice that any state which is not attainable from some qualitative state whose place is R3, by some path which includes some state whose place is R1, cannot represent a possible motion consistent with this data. The temporal relationship between the spatial measurements implies that in R3 and R1 the ball is always moving leftwards. FROB automatically makes these computations, and the possibilities which remain are shown in Fig. 10.

Suppose we also knew other facts about the ball, for example, that the ball was inelastic. Such facts can rule out particular states of motion (in this case, rebounding from a surface), and the same algorithms can combine these constraints with constraints from measurements. When inelasticity is assumed, for example, FROB concludes that the ball cannot ever be inside the well (see Fig. 11). While we do not know exactly what has happened, the pruning

algorithms built into FROB can still significantly constrain the possibilities.

V. DISCUSSION

This paper has presented a theory of interpreting measurements taken across time and has illustrated the utility of the theory through two extended examples. The problem it solves is central to qualitative physics since it provides a way to interpret information about the world in terms of qualitative physics. For the diagnosis problem this theory provides a general ability to test fault hypotheses to see if they actually explain the observed behavior.

Importantly, the theory relies on very few assumptions. As we have demonstrated, the small number of assumptions allows the theory to be applied to two very different representations and domains. The assumptions of a total envisionment (or the potential to generate it as needed) 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. No apparent reason exists why this theory cannot be used with device-centered models, such as [4] and [15], or discrete-process models, such as [14] and [16]. In fact, we expect that the constraints on partitioning numerical data will be the same for both device-centered and process-centered ontologies.

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, 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 capable of "reading gauges" for themselves could be automatically compiled from the structural description of a system.

Plans for Future Work

While exciting potential applications of this theory exist, there is still more research to do.

- The current implementation requires regular sequences as input. The gap-filling techniques which generate explicit interpretations for small gaps have not been implemented. We are extending the implementation to handle sparse data and fill small gaps.

- Currently, numerical information is translated into symbolic terms by hand. We are developing modules that automatically perform this translation for several domains, including motion and QP descriptions.
- We are formalizing the envisionment pruning algorithms developed for FROB to form a general-purpose *logic of occurrence*.
- We are developing a formal theory of measurement error and will use this theory to develop backtracking algorithms for handling noisy data. An important part of solving this problem is identifying domain-dependent constraints, such as continuity, which suggest how missing data should be filled in. Although some of these errors should be handled by the gap-filling techniques described earlier, we expect that, in general, some degree of backward communication with the initial signal-to-symbol processing will be required.
- We are developing heuristics, analogous to those of the "one-look" theory [9], to restrict possible interpretations still further. Introducing duration information, for example, allows the search through the envisionment for paths to fill gaps in sparse data to be cut off when the estimated duration of the path exceeds the duration of the gap.

ACKNOWLEDGMENT

Discussions with Johan de Kleer, Dan Weld, Brian Falkenhainer, and Dave Waltz led to significant improvements in both form and content. John Hogge provided invaluable programming assistance.

REFERENCES

- [1] J. Allen, "Maintaining knowledge about temporal intervals," Dept. Computer Science, Mass. Inst. Technol., Cambridge, Tech. Rep. TR-86, Jan. 1981.
- [2] D. Bobrow, Ed., *Qualitative Reasoning about Physical Systems*. Cambridge, MA: MIT Press, 1985.

- [3] D. M. Considine, Ed., *Chemical and Process Technology Encyclopedia*. New York: McGraw-Hill, 1974.
- [4] J. de Kleer and J. Brown, "A qualitative physics based on confluences," *Artificial Intelligence*, vol. 24, 1984.
- [5] K. Forbus, "Spatial and qualitative aspects of reasoning about motion," in *Proc. 1st Ann. Conf. American Association for Artificial Intelligence*, Aug. 1980.
- [6] —, "A study of qualitative and geometric knowledge in reasoning about motion," Mass. Inst. Technol. Artificial Intelligence Lab, Cambridge, MA, Tech. Rep. TR-615, Feb. 1981.
- [7] —, "Qualitative reasoning about physical processes," in *Proc. 7th Int. Joint Conf. Artificial Intelligence*, Aug. 1981.
- [8] —, "Qualitative reasoning about space and motion," in *Mental Models*, D. Gentner and A. Stevens, Eds. Hillsdale, NJ: Erlbaum, 1983.
- [9] —, "Measurement interpretation in qualitative process theory," in *Proc. IJCAI-8*, Karlsruhe, Germany, 1983.
- [10] —, "Qualitative process theory," *Artificial Intelligence*, vol. 24, pp. 85-168, 1984.
- [11] —, "Interpreting measurements in physical systems," in *Proc. AAAI-86*, Philadelphia, PA, Aug. 1986.
- [12] C. Hampel and G. Hawley, Eds., *The Encyclopedia of Chemistry*, 3rd ed. New York: Van Nostrand Reinhold, 1973.
- [13] D. R. Reddy, L. D. Erman, R. D. Fennel, and R. B. Newly, "The HEARSAY speech understanding system: An example of the recognition process," in *Proc. IJCAI-3*, 1973, pp. 185-193.
- [14] R. Simmons, "Representing and reasoning about change in geologic interpretation," Mass. Inst. Technol. Artificial Intelligence Lab, Cambridge, Tech. Rep. TR-749, Dec. 1983.
- [15] B. Williams, "Qualitative analysis of MOS circuits," *Artificial Intelligence*, vol. 24, 1984.
- [16] D. Weld, "Switching between discrete and continuous process models to predict genetic activity," Mass. Inst. Technol. Artificial Intelligence Lab, Cambridge, Tech. Rep. TR-793, Oct. 1984.



Kenneth D. Forbus received the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge.

He was associated with the MIT Artificial Intelligence Laboratory from 1973 to 1984. He is an Assistant Professor of Computer Science at University of Illinois at Urbana-Champaign. At Illinois he founded the Qualitative Reasoning Group, the first university research group in this area in the world. In addition to qualitative physics, his research interests include inference

engine design, truth-maintenance systems, engineering problem solving, cognitive simulation, and learning in physical domains.

Dr. Forbus is on the editorial board for the *International Journal for Artificial Intelligence in Engineering*, and is Science Program Chairman for the Sixth National Conference on Artificial Intelligence.