Qualitative State Spaces: A Formalization of the Naïve Physics Approach to Knowledge–Based Reasoning*

François E. Cellier and Nicolas Roddier
Department of Electrical and Computer Engineering
University of Arizona
Tucson, AZ 85721

Abstract

In this paper, we have attempted to formalize the Naïve Physics approach to knowledge–based qualitative reasoning in such a way that it can be compared with the quantitative analysis techniques used in system theory. It is shown that Naïve Physics models resemble classical quantitative models in more respects than is commonly assumed. This resemblance opens up an entire catalog of currently unanswered questions relating to Naïve Physics models. It helps prove some theorems about such qualitative models, but it also unveils some of the shortcomings of these models.

1 Introduction

Qualitative reasoning has become an important branch of A.I. research because it seems to play an important role in human decision making. The capability of automated decision making is important in particular in the context of Space exploration and colonization. In Space, human labor will be a scarce resource for a long time to come. Thus, we cannot conceive the conquest of our solar system without the deployment of *high autonomy systems*. In order to colonize our moon or the planet Mars, we must devise a technology in which autonomously operating robots can create a livable environment for humans. These robots will not be able to function properly unless they have the capability to make decisions on their own in a partially unknown environment.

How does the decision making process work? Humans make decisions by envisioning a set of possible scenarios (experiments), and by analyzing the effects of these scenarios using *mental models*. They "simulate" their mental models using *mental simulation*. The trajectories resulting from a mental simulation are referred to as *episodes*. They then go ahead with implementing the one scenario for which the mental simulation results in the most advantageous episode. If robots will ever be able to operate meaningfully on another planet, they must be able to mimic the human decision making process. They must be able to create new models of their environment on the fly, and use them immediately in "mental" simulation experiments. Thus, we need to design methodologies which allow us to model the human decision making process, i.e., we need to model the mechanisms of human understanding itself.

If I hold a glass full of water in my hand, and if I open my fingers, I *know* that the glass will fall down, and that, upon impact, it will break and spill the water all over my carpet. However, in my mental simulation, I don't solve any differential equations at all. I don't know *when* exactly the glass will hit the floor, and I don't know into how many pieces the glass will disintegrate, but often, these details are not important for proper decision making. The question thus is whether we can describe mental processes, that is: unprecise, qualitative models in precise mathematical terms.

The preceding example can be studied in the following way. I know that a gravitational force exists which has a tendency to pull objects downward. The precise amount of that gravitational force is often less important for decision making than knowledge of the direction in which it pulls objects. I know that if I grab the glass with my hands, it can't move. Since all objects which are exposed to a resulting force vector accelerate in the direction of that force vector, I know that there must exist a reaction force, thus keeping the glass in place. When I open my fingers, the gravitational force is no longer compensated for, and thus, the glass will accelerate in the direction of the gravitational force vector, i.e., the glass will begin to fall. After some time, the glass reaches the floor with

*Research supported by NASA-Ames Co-operative Agreement No. NCC 2-525, "A Simulation Environment for Laboratory Management by Robot Organizations."

0-8186-2162-1/91/0000/0040$01.00 © 1991 IEEE

a finite speed. I know that glasses don't penetrate floors. Thus, the glass needs to decelerate abruptly to a velocity of zero. An abrupt deceleration requires a very large (theoretically infinitely large) force to accomplish. I know that the internal structural stability of the glass cannot stand such a large force, and therefore, the glass will break. I finally know that broken glasses don't hold fluids very well, and thus, the water will be spilled over my carpet.

This example shows a knowledge-based approach to reasoning about the fate of my glass. Notice that the entire analysis was accomplished without mentioning any hard numbers. The gravitational force, the speed at impact, the force needed to decelerate the glass, etc. were nowhere quantified. No differential equation was solved to determine the trajectory behavior of my glass during its fall. All we needed in order to solve this problem were rudimentary bits and pieces of physical knowledge — this is what is referred to as "Naïve Physics." Knowledge-based qualitative reasoning refers to the set of methodologies which enable us to encode this type of imprecise knowledge in computer algorithms. Several researchers have tackled this problem. The most comprehensive summary of these efforts is given in [2].

However, it is not entirely clear that this is the way how humans reason about simple physical problems. Maybe, I never consider the gravitational force at all when I reason about the fate of my glass. Maybe, I know that the glass will fall down because I remember that my father, when I was 10 years old, let a two liter bottle of Chianti wine slip through his fingers, and our dog came, and licked the entire contents of the bottle from the carpet, and got himself completely drunk. But if this were true, how come that my brain correlated the water glass event with a seemingly unrelated event that happened more than three decades ago, taking into consideration that I observe and store an unbelievable manifold of different *episodes* every day of my life? How come a physician whom I had visited only once four years earlier introduced himself to me (since his office had misplaced my previous patient card), but frowned after the third word and exclaimed: "But you *have* been here before, haven't you?" although the guy must be seeing 50 patients a day? Our brain obviously has a remarkable capability to relate different, but seemingly similar, patterns to each other, and maybe, pattern recognition is at the heart of most "mental simulations."

Let me explain why I believe that pattern recognition is indeed a much more frequently used and much more powerful tool in assessing qualitatively the behavior of a system. I still own a dog who loves to play ball. I kick the ball with the side of my foot (I usually wear sandals, and a straight kick hurts my toes), and my dog runs after the ball as fast as he can. I was able to observe the following phenomenon: If I place my foot to the left of the ball, my dog will turn to the right to be able to run after the ball as soon as I hit it. He somehow *knows* that the ball will be kicked to the right. If I now change my strategy, and place my foot to the right of the ball, my dog immediately swings around to be ready to run to the left. He obviously has some primitive understanding of the mechanics involved in ball kicking. However, I assure you that I never let my dog near my physics texts, and thus, he had no opportunity to study Newton's laws — not even in their naïve form.

Thus, there exist two fundamentally different types of qualitative reasoners. One approach is *knowledge-based* and attempts to encode the *structure* of the system, while the other is *inductive* and attempts to encode the *behavior* of the system. Both approaches have their distinctive pro's and con's and deserve to be further pursued. In this paper, we shall concentrate on the knowledge-based approach, show what has been accomplished in this context, formalize the approach, bring it in relation with earlier work on quantitative models, and unravel its strengths and its weaknesses. A further elaboration of both approaches can be found in [5].

2 Definitions

Unfortunately, the literature on qualitative modeling in general and on Naïve Physics in particular is full of imprecisely defined, partially overlapping, and often even entirely redundant terminologies. Depending on the author, we meet terms such as *qualitative models*, *qualitative reasoning*, *qualitative physics*, *naïve physics*, and *common sense reasoning*. All of these terms are used in very similar contexts, and they are hardly ever properly defined. I shall use the following terminology:

1. *Qualitative variables* are variables which assume a finite ordered set of qualitative values, such as "minuscule," "small," "average," "large," and "gigantic." The literature on quantitative soft sciences is a little more precise on this definition than the literature on artificial intelligence. For instance, Babbie [1] distinguishes between:

   (a) *Nominal measures*, i.e., variables whose values have the only characteristics of exhaustiveness and mutual exclusiveness. Nominal measures are unordered sets. Typical nominal variables might be the religious affiliation, or the hair color of a person. Such variables are not useful as state–variables in a simulation. They can play a role as parameters.

   (b) *Ordinal measures*, i.e., variables who are nominal, and in addition, are rank–ordered. These variables are what I called above qualitative variables. However, sometimes we shall let go of the condition of mutual exclusiveness, for example, when we operate on *fuzzy sets.*

   (c) *Interval measures*, i.e., variables which are ordinal, and in addition, have the property that a distance measure can be defined between any two values, that is: interval variables can be added to and/or subtracted from each other. A typical candidate for a "soft" interval variable might be the intelligence quotient.

   (d) *Ratio measures*, i.e., variables which are interval measures, and in addition, have a true zero point.

2. *Qualitative behavior* denotes a time–ordered set of values of a qualitative variable, i.e., an episode.

3. *Qualitative models* are models that operate on qualitative states.

4. A *qualitative simulation* is an episode generator which infers qualitative behavior from a qualitative model.

Why are we interested in qualitative modeling and simulation? A number of applications for this methodology can be named:

1. *Incomplete System Knowledge.* Some details about the system under investigation are missing. Without such detail, quantitative simulation (at least in the sense of a single trajectory generation) cannot work. For example, after an anomaly is detected in a flight, the pilot will usually switch off the autopilot since s/he cannot trust any longer that the model which is an inherent part of the autopilot still reflects the behavior of the modified system adequately. Since the qualitative model operates on more highly aggregated variables, it may be somewhat more robust, i.e., it may be less sensitive to system modifications.

2. *Response to a Class of Experiments.* Until now, we always examined the response of a system to a single experiment. Sometimes it is more useful to examine the set of output trajectories resulting from applying an entire set of input trajectories to the system. Qualitative models are sometimes more adequate for this type of applications. Sensitivity analysis in the large is an alternative [3,4].

3. *Generalization for Decision Making.* It is sometimes hard to "see the forest for the trees." Quantitative models generate quantitative (that is: detailed) responses. It may be difficult to aggregate this detailed information in an automated system for purposes of knowledge generalization. It might be better not to generate this detailed knowledge in the first place. Qualitative modeling allows us to aggregate knowledge earlier in the game. This can sometimes be beneficial.

Evidently, a number of good reasons can be stated why we may wish to employ qualitative modeling and qualitative simulation. However, a number of bad reasons are also quite often heard:

1. "Qualitative simulation is cheaper than quantitative simulation. If quantitative simulation, in a real–time situation, cannot produce the results fast enough, qualitative simulation may be the answer to the problem." Wrong! Algorithms used for qualitative simulation are by no means faster than those used in quantitative simulation. In qualitative simulation, many alternative branches must generally be explored, whereas quantitative simulation usually produces one individual trajectory. Thus, quantitative simulation is normally faster than qualitative simulation if applicable. Thus, if your quantitative real–time simulation executes too slowly, don't go to qualitative simulation, go to a nearby computer store and buy yourself a faster computer.

2. "Qualitative simulation requires a less profound understanding of the mechanisms that we wish to simulate. Therefore, if we don't fully understand the mechanisms that we wish to simulate, quantitative simulation is out of the question, whereas qualitative simulation may still work." Wrong again! Qualitative simulation has as stringent constraints as quantitative simulation, they are just a little different. A convenient user interface relieves the user from some of the intricacies of detailed understanding of the simulation mechanisms, not the modeling methodology *per se.* Today's languages for quantitative simulation (such

as ACSL) are very user–friendly, more so than today's languages for qualitative simulation. This is due to the fact that quantitative simulation languages have been around much longer. Thus, if you don't understand what you are doing, don't go to qualitative simulation, go to an expert who does.

Qualitative modeling is not an *alternative* to quantitative modeling. When we have the knowledge available to produce a decent state–space model, it will in all likelihood work much better. Don't believe that, because *we* can't solve differential equations in our heads, our robots shouldn't do it either. Qualitative modeling presents us with an *enhancement* of our toolbox, and sometimes, this tool may be just the right one, but don't view (as unfortunately many of the researchers do) quantitative and qualitative modeling as in competition with each other. They are complementary rather than competitive techniques.

3 State Discretization and Landmarks

Naïve Physics models are characterized by a very limited set of qualitative values. Variables usually assume only one of three values: +, 0, and −. In order to maintain as much realism as possible within our highly aggregated state–space, we shall make the reasonable assumption that a continuous variable cannot jump from positive values to negative values without going through zero, and vice versa. Thus, we shall expand episodes as needed by adding values of 0 whenever the episode switches from + to −, or from − to +. + and − are two *regions*, whereas 0 is a *landmark*. In all Naïve Physics systems, adjacent regions are always separated by landmarks. The {− 0 +} set is the simplest of all meaningful sets of regions and landmarks.

Different authors use different approaches to alleviate the problem of reducing the true behavior of a physical system to a trivial behavior in the process of discretization. Kuipers [7] has a mechanism to invent new landmarks on the fly, and he includes the additional landmarks of $-\infty$ and $+\infty$ right from the beginning. Morgan [9,10] operates on the minimal set of regions and landmarks {− 0 +}, expanded only with two more elements: ? meaning *don't know*, and $\sqcup$ denoting an illegal or inconsistent value. ($\sqcup$ is a highly abstracted version of my Mac's trash can.) However, Morgan represents each variable through a *vector*:

$$x = (+\,+\,-) \tag{1}$$

meaning that x and $\dot{x}$ are currently positive, while $\ddot{x}$ is currently negative. Morgan carries along with the current value of a state variable information about its first and second time derivatives. Morgan's methodology can be easily modified to either include the first time derivative only, or to include higher time derivatives as well. We shall discuss the consequences of such a modification. However, for most applications, the above suggested triple seems to be optimal. In this paper, I shall basically follow the approach of Morgan since it comes closest to the classical concept of a state–space.

4 Operation on Qualitative Variables

Two qualitative variables x and y can easily be added. The truth table for adding qualitative variables is as follows:

$$ADD = \begin{array}{c} {}_x\backslash^y \\ - \\ 0 \\ + \\ ? \end{array} \begin{array}{c} \begin{array}{cccc} - & 0 & + & ? \end{array} \\ \begin{pmatrix} - & - & ? & ? \\ - & 0 & + & ? \\ ? & + & + & ? \\ ? & ? & ? & ? \end{pmatrix} \end{array} \tag{2}$$

Of course, x and y are in fact not scalars but *qualitative vectors.* However, we know that:

$$\begin{aligned} \text{if } z &= x + y \\ \rightarrow \dot{z} &= \dot{x} + \dot{y} \\ \rightarrow \ddot{z} &= \ddot{x} + \ddot{y} \end{aligned} \tag{3}$$

Similarly, we can define qualitative subtraction and multiplication operators with the following truth tables:

$$SUB = \begin{array}{c} {}_x\backslash^y \\ - \\ 0 \\ + \\ ? \end{array} \begin{array}{c} \begin{array}{cccc} - & 0 & + & ? \end{array} \\ \begin{pmatrix} ? & - & - & ? \\ + & 0 & - & ? \\ + & + & ? & ? \\ ? & ? & ? & ? \end{pmatrix} \end{array} \tag{4}$$

and:

$$MULT = \begin{array}{c} {}_x\backslash^y \\ - \\ 0 \\ + \\ ? \end{array} \begin{array}{c} \begin{array}{cccc} - & 0 & + & ? \end{array} \\ \begin{pmatrix} + & 0 & - & ? \\ 0 & 0 & 0 & 0 \\ - & 0 & + & ? \\ ? & 0 & ? & ? \end{pmatrix} \end{array} \tag{5}$$

In the multiplication, the corresponding vector function looks a little different since:

$$\begin{aligned} \text{if } z &= x \cdot y \\ \rightarrow \dot{z} &= x \cdot \dot{y} + \dot{x} \cdot y \\ \rightarrow \ddot{z} &= x \cdot \ddot{y} + 2 \cdot \dot{x} \cdot \dot{y} + \ddot{x} \cdot y \end{aligned} \tag{6}$$

The constant factor of 2 in the second derivative equation doesn't appear in the qualitative equations since $2 \cdot x$ is positive exactly if x is positive. Thus, multiplication of a qualitative variable with a positive constant is a *do nothing* operation.

The truth table for the qualitative division operator is not completely obvious. It has been defined as follows:

$$DIV = \begin{array}{c} {}_x\backslash^y \\ - \\ 0 \\ + \\ ? \end{array} \begin{array}{c} \begin{array}{cccc} - & 0 & + & ? \end{array} \\ \left(\begin{array}{cccc} + & \sqcup & - & ? \\ 0 & ? & 0 & 0 \\ - & \sqcup & + & ? \\ ? & \sqcup & ? & ? \end{array} \right) \end{array} \qquad (7)$$

5 Qualitative Simulation

Now, we are ready to perform a qualitative simulation. Let us explain by means of an example how this works. Figure 1 shows a simple RC circuit:

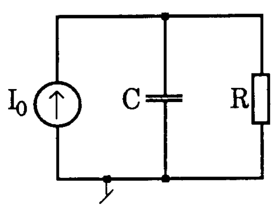

Figure 12.1. Simple RC circuit.

We wish to study the behavior of this circuit for a step function applied to the current source, assuming that the initial voltage across the capacitor is negative. This problem had originally been proposed by Williams [11].

A quantitative state–space description for this problem can be given as follows:

$$\dot{x} = \frac{-1}{R \cdot C} x + \frac{1}{C} u \qquad (8)$$

$$y = x \qquad (9)$$

where the input u is the current through the source, and the output y is the voltage across the capacitor. Using our software, QualSim, the qualitative version of this state–space description can be specified as follows:

```
// [xdot] = CQSTATE(x)
  u = CPLUS;
  xdot = QSUB(u, x);
RETURN
```

and:

```
// [y] = QOUT(x)
  y = x;
RETURN
```

where *CQSTATE* models the *continuous qualitative state equations*, and *QOUT* models the *qualitative output equations.* The qualitative subtraction (QSUB) of $u - x$ is the qualitative equivalent of the quantitative state equation since positive constant multipliers can be ignored in the qualitative model.

The initial conditions can be specified as:

```
// [x] = CQSTATEIC(dummy)
  x = IMINUS;
RETURN
```

The vector *IMINUS* denotes an *initial negative qualitative value*:

$$IMINUS = (\ -\ ?\ ?\ ) \qquad (10)$$

We happen to know that the initial capacitor voltage is negative, but we don't claim any knowledge about the initial value of the derivative of this voltage.

We can assign the initial value to x by calling *CQSTATEIC.* Thereafter, we can call *CQSTATE* to compute the initial value of the derivative. We find:

$$\begin{array}{ccc} t & x & \dot{x} \\ 0 & ((-\ ?\ ?) & (+\ ?\ ?)) \end{array}$$

However, since $x(2)$ must be equal to $\dot{x}(1)$, we can replace one of our question marks, and iterate once more:

$$\begin{array}{c} t \\ 0 \\ 0 \end{array} \begin{array}{c} \begin{array}{cc} x & \dot{x} \end{array} \\ \left(\begin{array}{cc} (-\ ?\ ?) & (+\ ?\ ?) \\ (-\ +\ ?) & (+\ -\ ?) \end{array} \right) \end{array}$$

We repeat the same process a second time:

$$\begin{array}{c} t \\ 0 \\ 0 \\ 0 \end{array} \begin{array}{c} \begin{array}{cc} x & \dot{x} \end{array} \\ \left(\begin{array}{cc} (-\ ?\ ?) & (+\ ?\ ?) \\ (-\ +\ ?) & (+\ -\ ?) \\ (-\ +\ -) & (+\ -\ +) \end{array} \right) \end{array}$$

and by now, all question marks have disappeared. The relation between x and $\dot{x}$ is called a *consistency constraint*, and the mechanism to ensure consistency among all such relations is called the process of *constraint propagation* [6].

Now, we are ready to *qualitatively integrate* the state vector to its next value, incrementing the *qualitative* *clock* to 1. Here, my methodology diverges from Morgan's. I use a *qualitative forward Euler* algorithm to integrate the state equations. The quantitative version of this algorithm can be written as:

$$x_{k+1} = x_k + \Delta t \cdot \dot{x}_k \tag{11}$$

and its qualitative counterpart can thus be written as:

$$x_{k+1} = QADD(x_k, \dot{x}_k) \tag{12}$$

which has been made available as:

$$x_{k+1} = QINT(x_k, \dot{x}_k) \tag{13}$$

QINT is a little more general than *QADD* since it can integrate an entire qualitative state vector whereas *QADD* operates only on one single qualitative variable (which is in itself a vector of length three). In our example, these two functions are identical since we analyze a first–order system.

The result of this integration is:

$$\begin{array}{c} t \\ 0 \\ 0 \\ 0 \\ 1 \end{array} \begin{array}{cc} x & \dot{x} \end{array} \\ \left(\begin{array}{cc} (-\ ?\ ?) & (+\ ?\ ?) \\ (-\ +\ ?) & (+\ -\ ?) \\ (-\ +\ -) & (+\ -\ +) \\ (?\ ?\ ?) & (?\ ?\ ?) \end{array}\right)$$

Unfortunately, our previous method of constraint propagation will not help us any further at this point.

We need to *expand* our search, i.e., replace the *?'s* by all possible combinations of −, 0, and +. However, here we can apply a new set of constraints which we call *continuity constraints*. The first continuity constraint had been mentioned previously. If x_k was +, x_{k+1} can only be + or 0, but never −, since no variable can jump from + to − without passing through 0 on the way. Of course, the same holds true for $\dot{x}_k$ and $\ddot{x}_k$. But more continuity constraints exist. For example, if x_k was + and $\dot{x}_k$ was either + or 0, then x_{k+1} must be +, and cannot be 0. The full set of continuity constraints is described in [5,9].

Repeated integration leads to:

$$\begin{array}{c|cc} t & x & \dot{x} \\ 0 & (-\ +\ -) & (+\ -\ +) \\ 1 & (0\ +\ -) & (+\ -\ +) \\ 2 & (+\ +\ -) & (+\ -\ +) \\ 3 & (+\ 0\ 0) & (0\ 0\ 0) \\ 4 & (+\ 0\ 0) & (0\ 0\ 0) \end{array}$$

After four steps, the system reaches a *continuous steady–state.*

QualSim provides a function:

$$[y, x] = CQSIM(nstp) \tag{14}$$

which computes the *continuous qualitative simulation* of the system described by the user coded functions *CQSTATE.CTR*, *CQSTATEIC.CTR*, and *QOUT.CTR* over *nstp* steps. The function:

$$QPLOT(y) \tag{15}$$

produces the graph shown in Figure 2.

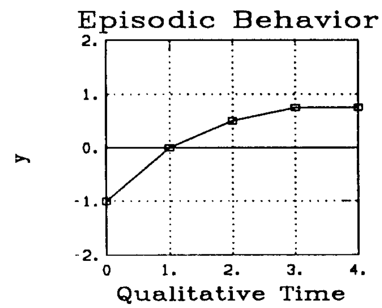

Figure 2. Episode of the RC circuit.

Figure 2 shows indeed the only possible physical solution to this problem. Encouraged by these nice results, let us try a somewhat more difficult problem:

$$\dot{x}_1 = +a \cdot x_2 \tag{16}$$

$$\dot{x}_2 = -b \cdot x_1 \tag{17}$$

which can be rewritten as:

$$\ddot{x}_1 = -(a \cdot b) \cdot x_1 = -\omega^2 \cdot x_1 \tag{18}$$

with the analytical solution:

$$x_1(t) = A \cdot \sin(\omega t) \tag{19}$$

$$x_2(t) = (A\omega) \cdot \cos(\omega t) \tag{20}$$

Let us check whether our qualitative simulation algorithm is able to reproduce this solution for us.

This problem is described in QualSim using the qualitative state–space model:

```
// [xdot] = CQSTATE(x)
  x1 = x(1 : 3);
  x2 = x(4 : 6);
  x1dot = x2;
  x2dot = QMINUS(x1);
  xdot = [x1dot, x2dot];
RETURN
```

It becomes evident that x is indeed the entire state vector which must first be unpacked. The qualitative output equations are coded as:

```
// [y] = QOUT(x)
  y = x;
RETURN
```

Since we wish to look at both state variables, we have no reason to unpack x in the first place. The initial conditions are assigned as follows:

```
// [x] = CQSTATEIC(dummy)
  x1 = IZERO;
  x2 = IPLUS;
  x = [x1, x2];
RETURN
```

We simulate this problem over eight integration steps, and find:

t	x_1	x_2	$\dot{x}_1$	$\dot{x}_2$
0	(0 + 0)	(+ 0 −)	(+ 0 −)	(0 − 0)
1	(+ + −)	(+ − −)	(+ − −)	(− − +)
2	(+ 0 −)	(0 − 0)	(0 − 0)	(− 0 +)
3	(+ − −)	(− − +)	(− − +)	(− + +)
4	(0 − 0)	(− 0 +)	(− 0 +)	(0 + 0)
5	(− − +)	(− + +)	(− + +)	(+ + −)
6	(− 0 +)	(0 + 0)	(0 + 0)	(+ 0 −)
7	(− + +)	(+ + −)	(+ + −)	(+ − −)
8	(0 + 0)	(+ 0 −)	(+ 0 −)	(0 − 0)

At this point, the system reaches a *periodic steady-state*. QualSim's qualitative plot function generates the following curves:

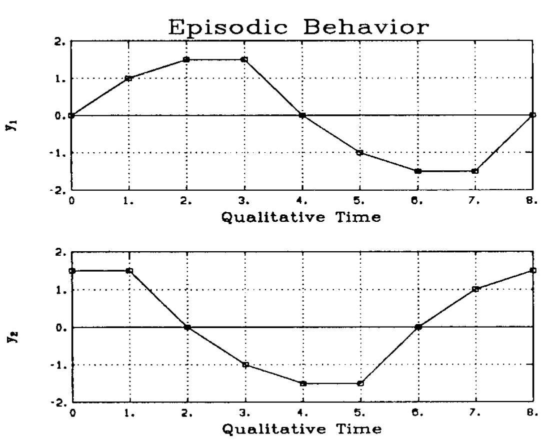

Figure 3. Episode of the sine wave model.

Unfortunately, the same approach will not work for more complex systems, such as a second order system with damping. This time, *CQSIM* is no longer able to eliminate all question marks by means of satisfying continuity constraints. Consequently, the episodical behavior splits into several branches which must all be explored, and very quickly, QualSim runs out of memory since there are simply too many branches to be followed.

Tony Morgan had an excellent idea how to overcome this problem [10]. Rather than trying to compute all episodes explicitly, he suggested to determine all legal qualitative states of the system, and then qualitatively integrate (or differentiate) over one step to find all legal successors of each of these legal states (which must, of course, also be among the legal states of the system).

Let me explain this concept by means of our second-order system with damping. In phase variables, this problem can be described as:

```
// [xdot] = CQSTATE(x)
  x1 = x(1 : 3);
  x2 = x(4 : 6);
  x1dot = x2;
  x2dot = QSUB(QMINUS(x1), x2);
  xdot = [x1dot, x2dot];
RETURN
```

We can evaluate all legal states by declaring the initial states as unknown. In this way, QualSim will determine the set of all feasible initial states which is, of course, identical with the set of legal states. In QualSim, this can be accomplished in the following way:

```
// [x] = CQSTATEIC(dummy)
  x1 = IWHAT;
  x2 = IWHAT;
  x = [x1, x2];
RETURN
```

The global variable *IWHAT* sets the qualitative vector to (? ? ?) (initial unknown). The QualSim statement:

$$[y, x] = CQSIM(1) \tag{21}$$

computes the set of legal states of this system, whereas the statement:

$$[y, x] = CQSIM(2) \tag{22}$$

computes the set of legal states and their immediate successors. The results of the analysis can be summarized as follows:

Table 1 State Transition Table of Second–Order System

state #	state	successor state #'s
(1)	$(- - +)$	(2)
(2)	$(- \ 0 \ +)$	(5)
(3)	$(- + -)$	(4), (7), (8)
(4)	$(- + \ 0)$	(3), (8)
(5)	$(- + +)$	(4)
(6)	$(0 \ - +)$	(1)
(7)	$(0 \ 0 \ 0)$	(7)
(8)	$(0 \ + -)$	(13)
(9)	$(+ - -)$	(10)
(10)	$(+ - \ 0)$	(6), (11)
(11)	$(+ - +)$	(6), (7), (10)
(12)	$(+ \ 0 \ -)$	(9)
(13)	$(+ + -)$	(12)

This is a so–called *finite state transition table.* While the information provided in the state transition table can be extracted from the x matrix generated by the *CQSIM* function, it is more convenient to use the statement:

$$fst = CQPERT(0) \qquad (23)$$

which calls *CQSIM*, and then postprocesses the data to generate the finite state transition table directly. *CQPERT* assumes that the system is specified using phase variables. *CQPERT* takes a dummy argument.

It is possible to represent the finite state transition table graphically in the form of a so–called *PERT network.* The PERT network of the above system is shown in Figure 4.

In a PERT network, each legal state is represented by a labeled node (a numbered circle). Transitions between states are indicated by directed paths. The two dashed paths in Figure 4 denote the soft transitions that lead to steady–state.

This method is indeed equivalent to exploring all episodes. Any given initial condition is represented by a legal state. Episodes starting from that initial condition can be found by simply following the directed paths through the PERT network. Whenever we pass through a node from which two or more transitions emanate, the qualitative behavior branches into several alternative episodes.

If the system is undercritically damped, its qualitative path will circle around the PERT network until, as $t \rightarrow \infty$, it finally moves into the steady–state node 7. If the system is overcritically damped, it will move to either state 3 or 11, whichever it reaches first, and stay in this node until, as $t \rightarrow \infty$, the system again moves into its steady–state node 7. Thus, the qualitative simulation has indeed been able to capture the physical behavior of this system correctly.

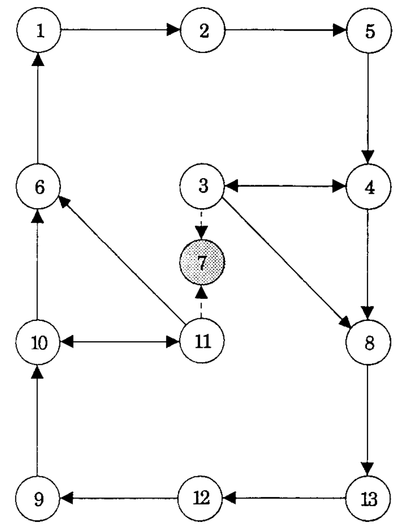

Figure 4. PERT network of second–order system with damping.

Unfortunately, as we proceed to a third order system, the PERT network becomes also unwieldy. The resulting state transition table contains 68 legal states, and up to eight paths emanate from a single node. The resulting PERT network is a mess. The graph is far from planar, and paths go in all directions. Thus, the best we can achieve here is to work with the finite state transition table directly.

6 Summary and Conclusions

Unfortunately, the current state–of–the–art Naïve Physics approach to modeling shows considerably more adversities than advantages. Let me list the problems that we are faced with:

1. The Naïve Physics approach did not live up to its promise to help us with the modeling process. We still need to derive a state–space model manually. Thus, the approach does not help us yet with automating the process of decision making. However, the approach saves us from the necessity of identifying model parameters through optimization which may potentially be a useful property.

2. While we were able to reduce the Naïve Physics model to the familiar state–space description (or rather its qualitative counterpart), the qualitative simulation has a tendency to explode, i.e.,

to branch out immediately into an unmanageable multitude of different episodes. Thus, qualitative simulation is not robust at all. We cannot blindly take any state–space model, convert it to its qualitative form, run QualSim, and hope to get anything meaningful out of it. However, robustness was exactly what we were after in the first place. Morgan's new approach using the finite state transition table helps alleviate this problem to some extent.

3. In order to get QualSim to work, we often require previous insight into the problem at hand, i.e., we must possess what we just try to gain — a most uncomfortable catch 22 situation.

4. Until now, QualSim can solve only trivial problems. The approach does not naturally extend to solving more complex problems involving higher order models. Thus, while QualSim allows us (in fact, forces us) to provide the system with *a priori* model information and while we are theoretically able to incorporate new knowledge on the fly (a frequently quoted deficiency of the inductive approach), we can in practice not do so, because the new knowledge will quickly enhance the system complexity to a degree at which QualSim fails to produce anything meaningful.

5. Many state–space models involve trigonometric functions such as *sin* or *cos*. No qualitative counterpart to those functions can be defined unless we can limit the range of their inputs (small signal behavior).

6. Naïve Physics allows us to specify parameters as being positive or negative only. This is not of much practical use. In most engineering problems, we know the parameter values with a tolerance of say $\pm 10\%$. Unfortunately, the Naïve Physics approach does not allow us to specify a limited range for a parameter. Consequently, the episodes produced by the qualitative simulation are usually too rich since they include the behavior of the system also for values of the parameters outside their physically meaningful range. Kuipers' QSIM program provides a partial answer to this problem[8]. QSIM operates on qualitative vectors of length two, but while the derivatives are always ternary variables of type $\{-\ 0\ +\}$, the state variables themselves may contain additional landmarks, and new landmarks can actually be discovered on the fly. This advantage was made possible (and is paid for) by forcing the user to formulate *all* constraints and the results of *all* qualitative operations explicitly, rather than relying on a set of implicit constraints and computational rules as QualSim does. This makes the problem formulation a little more difficult to derive for QSIM than for QualSim, but it makes QSIM a little more flexible than QualSim.

This paper has outlined the similarity between the quantitative state–space analysis of (especially linear) system theory and the knowledge–based approach to qualitative modeling. This similarity has its beauties and opens up a catalog of hitherto unanswered questions:

1. When we solve a quantitative linear control problem, we must always check for *controllability* and *observability* of our control system. However, problems with controllability and observability relate to singularities in the parameter space. Since our qualitative models refer to parameters only as ranges (e.g., p is positive) rather than as individual values, the concepts of controllability and observability have been thrown away in the transition. However in reality, the problem still exists. It would be useful to investigate whether there are unreachable qualitative states in a system, and under what conditions these exist. Using QualSim, the analysis is straightforward. We simply compute the state transition table of the system by letting both the initial states and the input assume values of *IWHAT*. However, we weren't able yet to come up with a general expression (similar to the controllability matrix of linear system theory) which would allow us to determine the existence of unreachable qualitative states conveniently and without resorting to exhaustive search.

2. In QualSim, we don't make the assumption of parameters being constant. If we let p be positive, p can be either constant or a function of time. Thus, while the model structure preserves the concept of *linearity*, the equally useful concept of *time-invariance* has been put to the sword in the transition. We might be able to get some of the power of that concept back by introducing a *qualitative similarity transformation.* Only those state transitions are physically valid which are generated by all similar qualitative state–space representations.

3. It is easy to define a qualitative equivalent to discrete–time systems. This has been shown in [5]. Many of the known properties of discrete-time systems can be reproduced in their qualitative counterparts. For example, it is known

that discrete–time first–order systems can oscillate while continuous–time first–order systems cannot oscillate. This behavior is reproduced in the qualitative analysis. Unfortunately, qualitative discrete–time models lead to even more ambiguity and have an even higher branching factor than qualitative continuous–time models. Yet, the concept is useful, and it would e.g. be interesting to check whether we can reproduce qualitatively the fact that non–linear discrete–time first–order models can lead to chaotic behavior (e.g., the logistic equation), while their continuous counterparts are never chaotic for any system order smaller than three.

Naïve Physics is still a wide open research field. While I am not yet convinced that this is really the right way to go, this is a fruitful area for research, since so many open questions remain to be answered. Only the future can tell whether this technique as a whole will survive, or whether it will end up in the ⊔ of science.

References

[1] Earl Babbie, *The Practice of Social Research,* fifth edition, Wadsworth Publishing Company, Belmont, CA, 1989.

[2] Daniel G. Bobrow, Ed., *Qualitative Reasoning About Physical Systems,* MIT Press, Cambridge, MA, 1985.

[3] François E. Cellier, "Enhanced Run–Time Experiments for Continuous System Simulation Languages," *Proceedings SCS MultiConference on Languages for Continuous System Simulation,* (F.E. Cellier, ed.), San Diego, CA, pp. 78–83, 1986.

[4] François E. Cellier and C. Magnus Rimvall, "Matrix Environments for Continuous System Modeling and Simulation," *Simulation,* **52**(4), pp. 141–149, 1989.

[5] François E. Cellier, *Continuous System Modeling,* Springer–Verlag, New York, 1991.

[6] Ernest Davis, "Constraint Propagation with Interval Labels," *Artificial Intelligence,* **32**, pp. 281–331, 1987.

[7] Benjamin Kuipers, "Qualitative Simulation," *Artificial Intelligence,* **29**, pp. 289–338, 1986.

[8] Benjamin Kuipers and Adam Farquhar, *QSIM: A Tool for Qualitative Simulation,* Internal Report: Artificial Intelligence Laboratory, The University of Texas at Austin, 1987.

[9] Antony J. Morgan, *The Qualitative Behaviour of Dynamic Physical Systems,* Ph.D. Dissertation, Wolfson College, The University of Cambridge, Cambridge, U.K., 1988.

[10] Antony J. Morgan, "Accuracy in Qualitative Descriptions of Behaviour," *Proceedings Winter Simulation Conference,* New Orleans, LA pp. 520–526, 1990.

[11] Brian C. Williams, "Qualitative Analysis of MOS Circuits," *Artificial Intelligence,* **24**, pp. 281–346, 1984.