Learning Qualitative Models for Systems with Multiple Operating Regions*

Sowmya Ramachandran, Raymond J. Mooney, Benjamin J. Kuipers
Department of Computer Sciences,
University of Texas,
Austin, Tx 78712

Email : sowmya@cs.utexas.edu

Abstract

The problem of learning qualitative models of physical systems from observations of its behaviour has been addressed by several researchers in recent years. Most current techniques limit themselves to learning a single qualitative differential equation to model the entire system. However, many systems have several qualitative differential equations underlying them. In this paper, we present an approach to learning the models for such systems. Our technique divides the behaviours into segments, each of which can be explained by a single qualitative differential equation. The qualitative model for each segment can be generated using any of the existing techniques for learning a single model. We show the results of applying our technique to several examples and demonstrate that it is effective.

Introduction

Qualitative reasoning is an elegant approach to studying the behaviour of a physical system without going into as much detail as in a numerical simulation. *Model building* and *model simulation* constitute the two major sub-problems of qualitative reasoning. There are several approaches to qualitative simulation, such as QSIM (Kuipers, 1986) and QPT (Forbus, 1984). Rapid advances have been made to improve the efficiency of the simulations and to fine tune them. However, the model building problem remains somewhat of an art form. Building models for a complex system requires significant knowledge of how the system works and is a time-consuming process.

Many researchers are addressing the problem of automatic model generation. One approach is to build models from existing libraries of model fragments (Forbus, 1984; deKleer and Brown, 1984; Crawford et al., 1990; Rickel, 1992). However, these techniques still require complete knowledge of all the model fragments. Another approach is to learn the model of a physical system from observations of its behaviour. Doyle (1988); Amsterdam (1993) have proposed techniques for learning models from behaviours using existing knowledge of processes and mechanisms commonly found in physical systems. These approaches are knowledge-intensive as well.

A number of researchers have formulated techniques for generating qualitative models of physical systems from a set of qualitative behaviours using inductive techniques (Coiera, 1989; Kraan et al., 1991; Richards et al., 1992; Dzeroski and Todorovski, 1993; Bratko et al., 1991). These require little knowledge of the system being modeled. Given a set of input behaviours, these techniques generate a single *qualitative differential equation (QDE)* that is consistent with the behaviours. The models generated are represented so that they can be used by QSIM.

Many complex physical systems cannot be described by a single QDE. They are explained by different QDEs that hold under different *operating conditions* or *regions*. For example, water boiling in a closed container requires three QDEs to explain its behaviour depending on whether it is below or at its boiling and whether all of the water has evaporated. A typical behaviour shown in Figure 1 passes through several of these regions.

However, MISQ (Kraan et al., 1991; Richards et al., 1992) and the other systems cannot learn models described by multiple QDEs. Given the behaviour in Figure 2, they will incorrectly learn a single QDE that explains the entire behaviour. It is desirable to have inductive model generators that can recognise that the behaviour is best described by multiple QDEs and learn them.

This paper describes a simple technique for automatically recognising that there are multiple QDEs underlying a physical system. It assumes only the QSIM (Kuipers, 1986) formalism and is independent of the induction algorithm used to generate

*This work has taken place in the Qualitative Reasoning Group at the Artificial Intelligence Laboratory, The University of Texas at Austin. Reason of the Qualitative Reasoning Group is supported in part by NSF grants IRI-8904454, IRI-9017047, and IRI-9216584, and by NASA contracts NCC 2-760 and NAG 9-665.

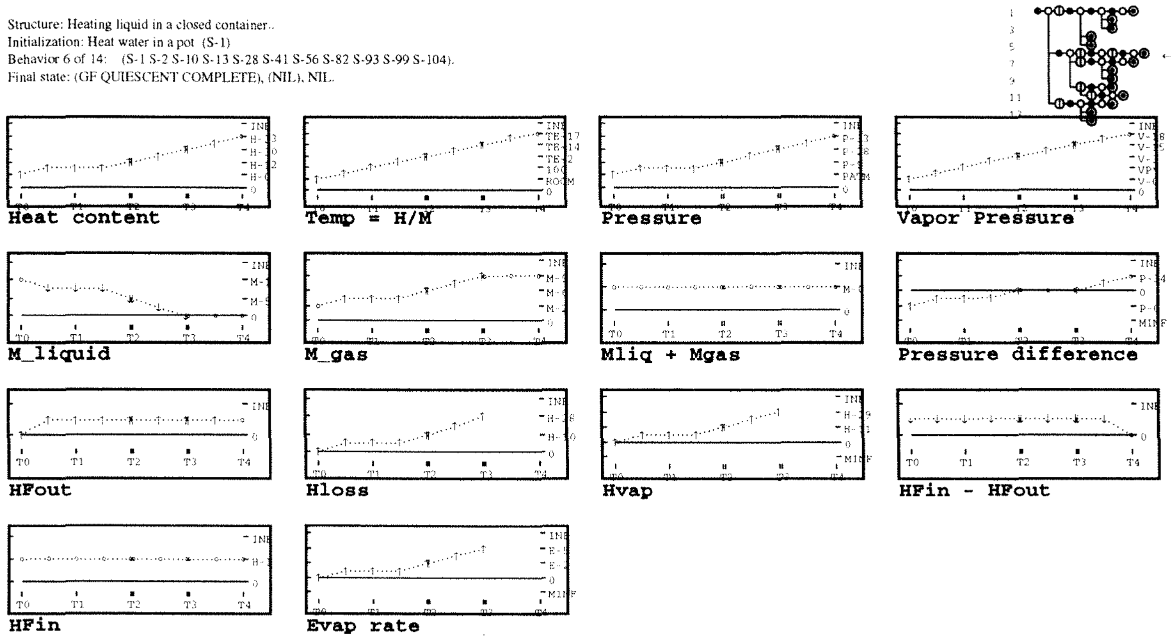

Figure 1: The Boiling-Water System: Typical Behaviour

the model. Given qualitative observations of the behaviours of a system, this technique identifies the various operating regions of the system where different QDEs hold.

We have evaluated our technique on the following physical systems: the Plant Water-balance system, the Boiling-Water system and the Divided-Tank system. Ultimately, we would like to be able to learn multiple QDEs from quantitative data as well.

In the next section, we will describe our technique in detail. Section 3 describes an experimental evaluation of this technique. In Section 4, we discuss future directions.

Learning Models with Multiple QDEs

A QDE is valid over some operating region. The conditions, expressed in terms of the values of the variables involved, over which the QDE is valid are called the *operating conditions* of the QDE. The movement of a system from the operating region of one QDE to that of another is called a *region transition.*

Behaviours of physical systems that pass through multiple operating regions exhibit region transitions. The segment between two consecutive transitions is governed by a single QDE. Our approach to learning models with multiple QDEs is to first break up the behaviours into such segments. Then, the system can use any of the existing induction algorithms to generate the QDE for each segment.

Thus, the problem of learning models for systems with multiple operating regions can be divided into the following sub-problems.

1. Break up the example behaviours into segments that fall within a single operating region.
2. Learn the QDE for each of the segments.
3. Identify the operating conditions for each QDE.
4. Unify the QDEs that describe the same operating region.

Although step 2 involves an induction algorithm to learn QDEs for each region, in the following subsections we describe techniques for performing steps 1, 3 and 4 that are independent of the learning algorithm used.

Step 1: Breaking up the behaviours into segments

To break up a behaviour into segments, it is sufficient to detect the time points where the behaviour moves from one region to another, i.e., to detect the region transitions. Our system uses the following heuristics to recognise transition points.

- **Discontinuous-Change heuristic**

  One of the interpretations of a region transition is that the actual state of the mechanism undergoes a discontinuous change (Kuipers, 1994). Thus, a discontinuity in the behaviour of any variable in the system can be used to detect a region transition. A discontinuity can be any of the following kind.

  1. Discontinuity in the magnitude of a variable. An example of this is a variable whose magnitude goes from being positive to being negative without going through zero.

     Figure 2 shows the variable *uptake* undergoing a discontinuous change in magnitude at time point T2.

  2. Discontinuity in the sign of the derivative of a variable. For instance, this happens when the derivative of a variable goes from being positive to being negative without going through zero.

     For example, in Figure 2, the derivative of the variable *net inflow* undergoes a discontinuous change at time point T2.

  This heuristic is justified by the fact that QSIM does not predict discontinuous behaviour unless it encounters a region transition that introduces the discontinuity. In order to learn a model that covers the given discontinuous behaviour, MISQ has to hypothesise a region transition at the point of discontinuity.

- **Non-analytic-Function heuristic**

  This heuristic relies on the properties of a certain class of functions called *analytic* functions. If a function is analytic over an interval, and is constant over any open sub-interval, it must be constant over the entire interval (Kuipers, 1994). Thus, under the assumption that the actual behaviours exhibited by all the variables in the system being modeled are analytic, if a variable is observed to be constant over some interval, but not over some other interval in the same behaviour, then the two intervals must be governed by different constraints and hence different QDEs.

  For example, in the behaviour shown in figure 2, the variable *turgor* exhibits non-analytic behaviour. It is non-constant over the interval (T0, T5) and is constant over the interval (T5,T6).

  This heuristic is justified because, under the *analytic function assumption*, QSIM will never generate a non-analytic behaviour unless it encounters a region transition that introduces the non-analyticity.

Step 2: Learning the QSIM models for each operating region

The previous section described the heuristics to break up behaviours into segments corresponding to different operating regions. The learner can now generate a single QDE to model each segment using any of the existing induction algorithms (Coiera, 1989; Kraan et al., 1991; Richards et al., 1992; Dzeroski and Todorovski, 1993; Bratko et al., 1991).

In our implementation, we have used MISQ (Kraan et al., 1991; Richards et al., 1992) to generate the QDEs. Given a set of qualitative behaviours, MISQ uses a *most specific generalisation* algorithm to generate QSIM models that are guaranteed to be consistent with the behaviours.

Step 3: Identifying the operating conditions for each region

Identifying the operating condition for each QDE can be thought of as an inductive process. The behaviour segment associated with each QDE provides positive examples of the condition under which it is active. We have used a most specific conjunctive generalisation approach to induce the operating conditions from positive examples, where the operating condition induced for a region is the range of all the qualitative values observed for each variable in the behaviour segment. Table 1 shows the operating condition for the region between T0 and T1 in the behaviour shown in Figure 2.

An operating condition represents the interior of a region, whereas QSIM represents a region by its boundaries. The region of validity of each QDE is specified in terms of transition mappings. These mappings specify *transition conditions*, i.e., conditions under which the system moves out of the operating region of one QDE into that of another.

It is easy to derive the boundary conditions of each QDE from its operating conditions. We view the variables that define the boundary of a QDE as *triggers* that cause the transition out of the region. Table 4 shows some examples of triggers that cause transitions between regions. The following observations help us in identifying such triggers.

1. Since discontinuities cannot occur spontaneously, a variable that changes discontinuously cannot cause a transition.
2. A variable that is steady over a region or does not cross a landmark value cannot cause a transition.

The rest of the variables are potential triggers and the transition condition is specified as the conjunction of their boundary values.

Step 4: Unification of regions

At the end of step 3, there are as many QDEs as there are behaviour segments. However, many of the segments could have the same underlying

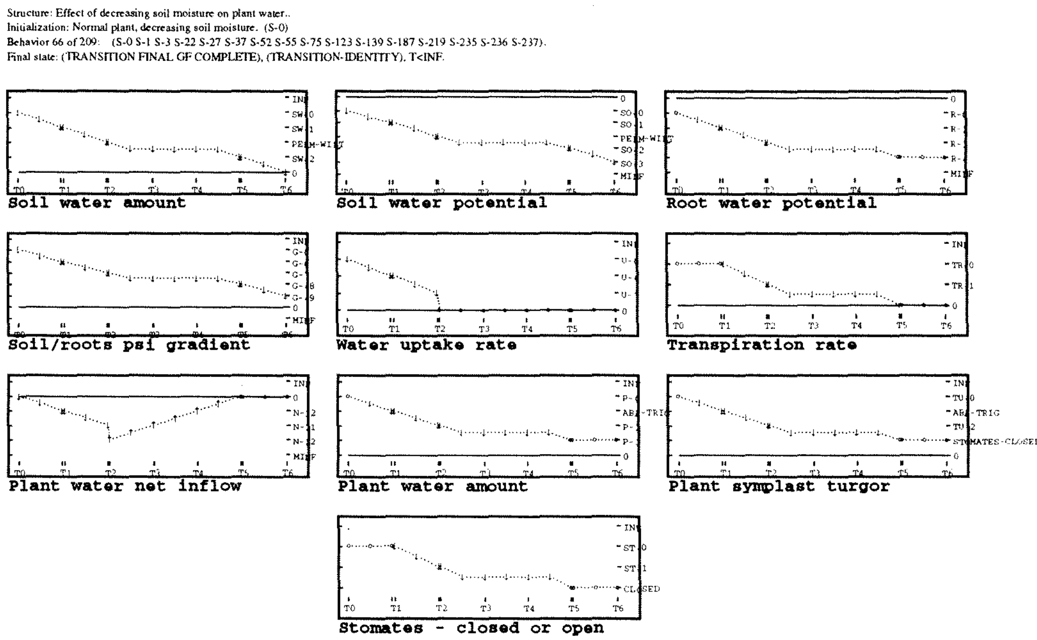

Figure 2: The Plant Water-Balance System: Input Behaviour

Variable	Qualitative interval
soil water amount	(SW-0 SW-1)
soil water potential	(SO-0 SO-1)
root water potential	(R-0 R-2)
soil/root psi gradient	(G-0 G-2)
water uptake rate	(U-0 U-6)
transpiration rate	TR-0
plant water net inflow	(0 N-12)
plant water amount	(P-0 ABA-TRIG)
plant symplast turgor	(TU-0 ABA-TRIG)
stomates	ST-0

Table 1: An Example of an operating condition

model. It is important to identify such segments and unify them.

A simple criterion to decide whether a set of regions should be unified is to check the set of constraints in the QDEs associated with the regions. In principle, two QDEs are identical if and only if their constraints sets are identical. However, since each QDE is learned inductively, two QDEs that should be identical may have different constraint sets. In practice, we have to rely on heuristics to guide region unification. We use two heuristics to identify the regions to be unified.

1. **Identical Constraints heuristic**
   If two regions have QDEs with identical sets of constraints, then they are unified. The operating condition for the unified region is the disjunction of the operating conditions of the individual regions, if they are *disjoint*. Otherwise, the two operating regions are combined by combining the qualitative intervals for each variable across the regions unified. Two operating conditions are *disjoint* if and only if there is at least one variable with non-overlapping qualitative values across the two regions.
2. **Identical Operating Conditions heuristic**
   Two regions are unified if they have identical operating conditions. The set of constraints defining the unified QDE is the intersection of the sets of constraints in the regions unified.

Our system applies the two heuristics repeatedly until no more regions can be unified.

Experimental Evaluation

We have implemented this technique in a system called MISQ-RT. We have used MISQ-RT to generate multiple QDE models for the following physical systems: the Plant Water-balance system, the Boiling-Water system and the Divided-Tank system.

We will first outline our experimental methodology and then describe the results for each of the three systems.

Methodology

We designed our experiments to test the effectiveness of our approach and our heuristics. We considered the heuristics for the identification of region transitions to have been effective if they could successfully detect all the region transitions. We evaluated the heuristics for region unification by comparing the regions generated by MISQ-RT with a model of the same physical system generated by an expert. We expected the heuristics to unify exactly those regions deemed identical by the expert. The specification of the conditions for region transitions were also evaluated by comparing the generated models with the expert model.

In each of the experiments, we generated behaviours by using QSIM to simulate the expert model. We picked a few of these behaviours as input for MISQ-RT, making sure that the selected behaviours traced different trajectories through the regions and exhibited different region transitions. The input to MISQ-RT also included totally ordered quantity spaces and dimensions for each variable.

The output from MISQ-RT was the generated model in a format that could be used by QSIM. The specification for each QDE included the constraints and the transition mappings.

QSIM indicates region transitions in the behaviours that it generates. This made it easy to check if MISQ-RT has identified all the region transitions in a behaviour.

Since the behaviours were generated by QSIM, there was a correspondence between the behaviors and the QDEs in the expert model that generated them. Thus, we could establish a correspondence between the regions generated by MISQ-RT and the expert QDEs. This was crucial in evaluating the heuristics for region unification. We expected MISQ-RT to unify only those regions that correspond to the same expert QDE.

The following subsections present the results of our experiments, followed by a discussion of the results.

Results

Water Boiling in a Closed Container This experiment modeled the scenario of water being brought to boil in a closed container. The model defined by the expert had *three* QDEs, *Heating*, *Boiling* and *Gas-only*. The QDE *Heating* is active when the water is being heated up to its boiling point. *Boiling* is active when the water has reached its boiling point. *Gas-only* is active when all the water has evaporated.

Figures 1 shows one of the four inputs to MISQ-RT. The other input behaviours showed all the water evaporating before it reached its boiling point. MISQ-RT identified *three* regions as well.

Before region unification, MISQ-RT generated as many QDEs as the number of behaviour segments between transitions. In this case, it generated *nine* QDEs. Table 6 shows the correspondence between these QDEs and those defined by the expert. Table 7 shows the correspondence between the generated QDE and the expert QDE after unification. MISQ-RT identified exactly those regions corresponding to the same QDE in the expert model.

Table 4 shows the transitions for each of the QDE as defined by the expert. Table 5 shows the transitions learned by MISQ-RT for each QDE. These tables show that MISQ-RT was successful in identifying all the variables that cause transitions out of each QDE. The transition condition for the

Generated QDE	Expert QDE
r106498	Boiling
r106493	Heating
r106487	Heating
r106481	Heating
r106464	Heating
r106486	Gas-only
r106479	Gas-only
r106492	Gas-only
r106503	Gas-only

Table 2: The Boiling-Water System: Region correspondence before unification

Unified QDE	Component QDE	Expert QDE
q106521	r106498	Boiling
q106522	r106493 r106487 r106481 r106464	Heating
q106523	r106486 r106479 r106492 r106503	Gas-only

Table 3: The Boiling-Water System: Region correspondence after unification

QDE *q106521* seems overly-specific. However, it is merely redundant since the two conditions would occur simultaneously. This brings up the question of the interpretation of the landmark value M-9 for the variable *mgas*. This is discussed in the section 4.

Plant Water-balance system In this experiment, we modeled a plant balancing the amount of water in its system, as the level of water in surrounding soil decreases. The model defined *four* QDEs, *healthy-stomates-closed-uptake*, *water-stress-uptake*, *healthy-stomates-closed-no-uptake*, *water-stress-no-uptake*. The main factors that determine the region that is active at any point are (1) whether the plant is healthy, and (2) whether there is any uptake of water from the soil. When the plant is *healthy*, i.e., when the concentration of water in its system is above a threshold, the size of stomatal-opening is constant. When the plant is *water-stressed*, i.e., not healthy, the stomates start closing. When the concentration of water in the soil falls below a certain level, there is no uptake of water. When the *stomates* are closed, the size of the stomatal-opening is constant. Thus, both the *healthy* and the *stomates-closed* condition result in the same QDEs. This example is similar to the one described in (Rickel and Porter, 1992).

Figures 2 shows one of the two input behaviours to MISQ-RT. MISQ-RT identified *six* QDEs, whereas the expert model had only *four*.

Table 6 shows the correspondence between the generated QDEs and the expert QDEs. Table 7 shows the correspondence after unification.

MISQ-RT unified only those QDEs that corresponded to the same QDE in the expert model. However, it did not unify all such QDEs. It could not identify that the QDE for the condition of a plant being *healthy* with the QDE for the condition when the *stomates* are *closed*. In region *q105961*, the plant is *healthy* but the *stomates* are open. Whereas in region *q105963*, the plant is *unhealthy* but the *stomates* are closed. These two conditions are explained by the same QDE in the expert model. However, this would affect neither the correctness of the model nor its generality.

Table 8 shows the transitions for each of the four QDEs defined by the expert. Table 9 shows the transitions for the QDEs generated by MISQ-RT. The transition conditions for the generated QDEs are more specific than for the expert QDEs. MISQ-RT did not propose transitions that did not occur in the input behaviours. For example, for the QDE corresponding to the region *water-stress-uptake*, it did not generate a transition when the variable *turgor* increases beyond the value *aba-trig*. This is because it did not encounter any behaviour with such a transition. We are investigating techniques for proposing transitions that did not occur in the input behaviours. This can be done by examining the generated regions pairwise for variables that could cause transitions between them.

The Divided-Tank system Figure 3 shows the Divided-Tank system. It has a tank with a partition in the middle. There is an inflow into region A and a drain in each of the regions A and B. This example is from (Soderman and Stromberg, 1991).

Figure 4 shows the behaviour tree and one of the predicted behaviours when the tank is filled from empty. QSIM predicted 7 behaviours, all of which were included in the input to MISQ-RT.

The expert model defined three QDEs, *FillA*,

From QDE	Transition Condition Variable	Qmag	Qdir	To QDE
Heating	pdiff	0	inc	Boiling
	Mliq	0	dec	Gas-only
Boiling	Mliq	0	dec	Gas-only
Gas-only				No transition

Table 4: The Boiling-Water System: Transition table for the expert model

From QDE	Transition Condition Variable	Qmag	Qdir	To QDE
q106522	pdiff	0	inc	q106521
	Mliq Mgas	0 M-9	dec NIL	q106523
q106521	Mliq Mgas	0 M-9	dec NIL	q106523
q106523				No transition

Table 5: The Boiling-Water System: Transition table for the generated model

Generated QDE	Expert QDE
r105904	water-stress-uptake
r105918	water-stress-no-uptake
r105924	healthy-stomates-closed-uptake
r105899	healthy-stomates-closed-uptake
r105929	healthy-stomates-closed-no-uptake
r105934	water-stress-no-uptake
r105939	healthy-stomates-closed-no-uptake
r105923	healthy-stomates-closed-no-uptake

Table 6: The Plant Water-Balance System: Region correspondence before unification

Unified QDE	Component QDEs	Expert QDE
q105958	r105904	water-stress-uptake
q105959	r105918	water-stress-no-uptake
q105960	r105924 r105899	healthy-stomates-closed-uptake
q105961	r105929	healthy-stomates-closed-no-uptake
q105962	r105934	water-stress-no-uptake
q105963	r105939 r105923	healthy-stomates-closed-no-uptake

Table 7: The Plant Water-Balance System: Region correspondence after unification

From QDE	Transition Condition Variable	Qmag	Qdir	To QDE
healthy-stomates-closed-uptake	turgor	aba-trig	dec	water-stress-uptake
	soil-psi	perm-wilt	dec	healthy-stomates-closed-no-uptake
water-stress-uptake	turgor	aba-trig	inc	healthy-stomates-closed-uptake
	stomates	closed	nil	healthy-stomates-closed-uptake
	soil-psi	perm-wilt	dec	water-stress-no-uptake
healthy-stomates-closed-no-uptake	turgor	aba-trig	dec	water-stress-no-uptake
	soil-psi	perm-wilt	inc	healthy-stomates-closed-uptake
water-stress-no-uptake	turgor	aba-trig	inc	healthy-stomates-closed-no-uptake
	stomates	closed	nil	healthy-stomates-closed-no-uptake
	soil-psi	perm-wilt	inc	water-stress-uptake

Table 8: The Plant Water-Balance System: Transition table for the expert model

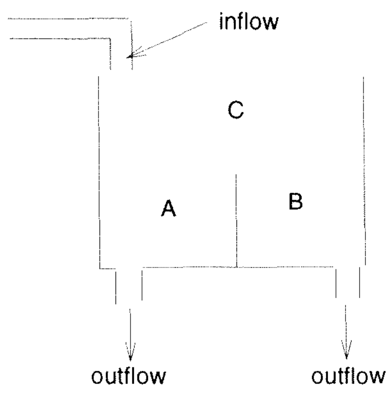

Figure 3: The Divided Tank system

From QDE	Transition Condition			To QDE
	Variable	**Qmag**	**Qdir**	
q105960	pwater	aba-trig	dec	q10598
	turgor	aba-trig	dec	
	swater	perm-wilt	dec	q105961
	soil-psi	perm-wilt	dec	
	uptake	0	dec	
	netflow	n-9	dec	
q105958	swater	perm-wilt	dec	q105959
	soil-psi	perm-wilt	dec	
	uptake	0	dec	
q105961	pwater	aba-trig	dec	q105962
	turgor	aba-trig	dec	
q105963				No Transitions
q105959	root-psi	R-4	dec	q105963
	pwater	P-3	dec	
	transp	0	dec	
	turgor	stomates-closed	dec	
	netflow	0	inc	
v	stomates	closed	dec	
q105962	root-psi	r-4	dec	q105963
	pwater	p-3	dec	
	transp	0	dec	
	turgor	stomates-closed	dec	
	netflow	0	inc	
	stomates-closed	closed	dec	

Table 9: The Plant Water-Balance System: Transition table for the generated model

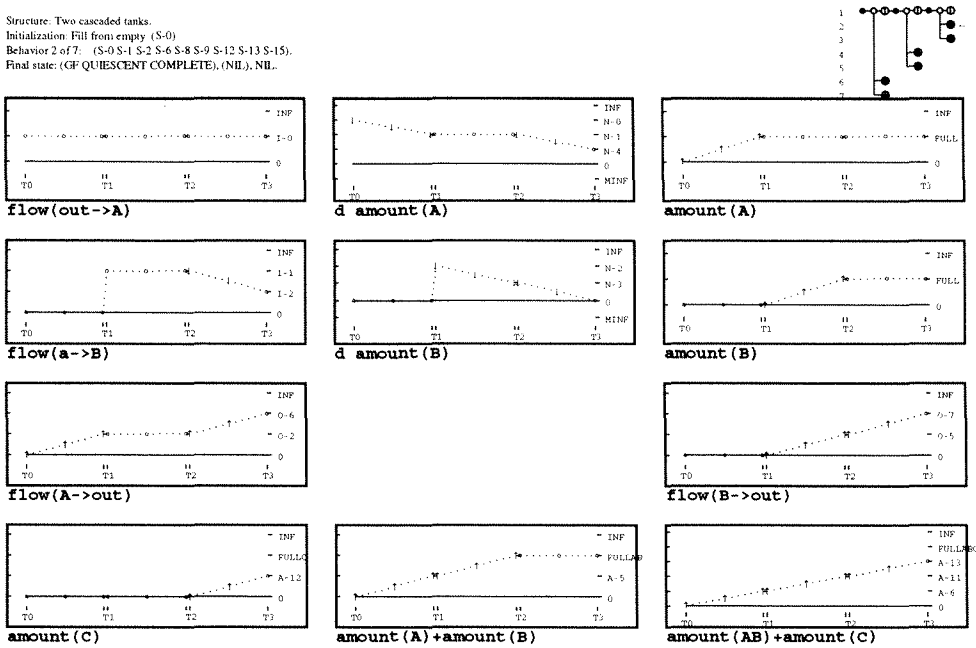

Figure 4: The Divided-Tank System: Input Behaviour

FillB and *Fill-both*. FillA is active when region A is being filled. *FillB* is active when region A is full and region B is getting filled. *Fill-both* is active when both A and B are full and region C is getting filled.

The model generated by MISQ-RT defined the same number of QDEs as the expert model. It unified exactly those QDEs that corresponded to the same QDE in the expert model. It identified all the variables responsible for causing transitions out of a QDE. Some of the transition conditions were overly-specific.

Discussion

The outcome of the experiments showed that the heuristics for identifying region transitions in qualitative behaviours were effective. In all of the experiments, MISQ-RT identified all the region transitions. These heuristics are independent of the induction algorithm used to learn the QDEs.

The heuristics for recognising and unifying identical regions were effective in unifying a large proportion of the QDEs deemed identical by the expert. They never made the mistake of unifying regions that the expert did not consider identical.

Sometimes, the technique generated overly specific transition conditions. This could cause QSIM to miss some transitions when it uses the generated model for simulation. We are investigating techniques for avoiding this through the use of negative examples of transitions conditions, i.e., situations where a proposed transition condition did not lead to a transition.

Since we used MISQ as our induction module, the QDEs generated for each region was guaranteed to be consistent with the behaviour segments for each region.

Future Work

Evaluation on more complex systems

The techniques we have proposed here rely on heuristics. These heuristics have to be validated by extensive experiments. We would like to perform experiments on systems more complex than those we have studied so far. One such system is the Reaction Control System (RCS) of the space shuttle (Kay, 1992).

Although, it is not desirable to learn a single QDE to explain the behaviour of a system with multiple operating regions, generating too many QDEs would adversely affect the generality of the model. The number of regions generated should be of an order less than the number of behaviour segments in the input. We would like to study the number of behaviours generated by MISQ-RT in relation to the number of behaviour segments in the input. We would also like to evaluate the generality of the models generated by our technique by using QSIM to simulate them and see how successful they are in predicting behaviours previously unseen.

Identifying Region Transitions from Quantitative Data

Our current implementation of the technique requires qualitative behaviours as input. However, to be useful in modeling real systems like the RCS, the system should be able to handle quantitative data as well.

The heuristics for identifying region transitions from behaviours can be applied to quantitative data as well. Non-analytic behaviour and discontinuous changes will have to be detected from quantitative data. So far we have not investigated techniques for doing this. This task is further complicated in the presence of noise. We plan to address this issue in the near future.

Matching Landmarks across Behaviours

The experiments have shown the *identical operating conditions heuristic* to be quite useful in unifying regions. However, the criterion for matching the operating conditions for regions is purely syntactic. Two operating conditions are considered identical if the qualitative intervals for each of the variables are identical across the two regions. Two qualitative intervals are considered identical is they are bounded by the same landmark value.

When the inputs to the learner come from different sources, landmarks may not match syntactically, even if they stand for the same event. Consider the situation where MISQ-RT is modeling a bathtub and it receives behaviours from two bathtubs, A and B. Let *FullA* and *Fullb* be the landmarks for the the event when A and B are full to capacity, respectively. As MISQ-RT processes the quantitative data, it has to recognise that *FullA* and *FullB* are qualitatively the same landmarks values, though they have different quantitative values. If they are not recognised to be the same, the operating conditions for the QDEs proposed for the two bathtubs will not match.

There are certain landmarks like *zero* that are special and can be matched easily across behaviours. A possible approach to this problem would be to use these special landmarks and the qualitative trends in the behaviour to match the other landmarks. This is an interesting problem that could arise in other applications of Qualitative Reasoning and is worth pursuing.

Related Work

Many techniques have been proposed for learning models of physical systems from observations of their behaviour (Coiera, 1989; Kraan et al., 1991; Richards et al., 1992; Dzeroski and Todorovski, 1993; Bratko et al., 1991). All of these, however, can only generate a single QDE model.

Falkenhainer (1990) describes a technique for building models for systems by analogy with other systems. This approach requires knowledge in the form of a library of processes. This is also true of (Rickel, 1992; Rickel and Porter, 1992) who describe a method for automatically building models from process libraries. Since they work at the process level, they are not concerned with region transitions directly.

The machine discovery system ABACUS (Falkenhainer and Michalski, 1986) learns the mathematical equations describing a set of numerical data. It can discover multiple equations that apply under different conditions. Although their technique does learn qualitative relations between variables, its main focus is on learning quantitative laws. The quantitative laws help in recognising and learning the various operating regions of the system. MISQ-RT, on the other hand, uses qualitative heuristics to identify the different operating regions. Thus, it can be used with quantitative as well as qualitative data.

Soderman and Stromberg (1991) have proposed a technique for learning models of systems that abruptly change between linear modes of operation. They address similar issues such as identifying "jump" in behaviours, finding correspondences between various segments of the behaviour and finding conditions under which each mode is active. They use system identification techniques to detect region transitions and to fit a model for each segment. However, they have to specify the model structure in advance. Our approach does not make any assumptions about the model structure. They also require knowledge in the form of bond graphs. The models they fit are quantitative models. Our approach works on qualitative behaviours and can be used in situations where quantitative observations are not available.

Nordhausen and Langley (1993) have proposed a technique for empirically discovering the laws that govern scientific phenomena. Their system discovers both qualitative and quantitative laws. It can also identify and model the different operating regions of the system. However, the system does not break up the observations into different segments automatically. This information has to be given as input to the system. The transition conditions have also to be specified with the input.

Conclusion

In this paper, we have proposed a method for learning models with multiple QDEs from qualitative observations of their behaviour. We have proposed heuristics to detect region transitions and for identifying corresponding regions. We have also suggested a technique for identifying the operating conditions of each QDE. The experiments reported here indicate that our approach is effective in identifying region transitions and learning models with multiple QDEs.

Acknowledgment

We thank James Lester and Raman Rajagopalan for comments on earlier versions of the paper. We

would also like to thank Bob Schrag for helping with the figures.

References

Amsterdam, J. (1993). Automated qualitative modeling of dynamic physical systems. Technical Report 1412, MIT Artificial Intelligence Laboratory.

Bratko, I., Muggleton, S., and Varsek, A. (1991). Learning qualitative models of dynamic systems. In *Proceedings of the Eighth International Workshop on Machine Learning.*

Coiera, E. (1989). Generating qualitative models from example behaviours. Technical Report 8901, University of New South Wales.

Crawford, J., Farquhar, A., and Kuipers, B. (1990). QPC: A compiler from physical models into qualitative differential equations. In *Proceedings of the Eighth National Conference on Artificial Intelligence.*

deKleer, J., and Brown, J. (1984). A qualitative physics based on confluences. *Artificial Intelligence*, 24:7–83.

Doyle, R. (1988). *Hypothesizing Device Mechanisms: Opening Up the Black Box.* PhD thesis, Massachusetts Institute of Technology.

Dzeroski, S., and Todorovski, L. (1993). Discovering dynamics: From inductive logic programming to machine discovery. In *Knowledge Discovery in Databases Workshop.*

Falkenhainer, B. (1990). A unified approach to explanation and theory formation. In Shrager, J., and Langley, P., editors, *Computational Models of Scientific Discovery and Theory Formation*, 157–196. Morgan Kaufmann.

Falkenhainer, B. C., and Michalski, R. (1986). Integrating quantitative and qualitative discovery: The ABACUS system. *Machine Learning*, 1:367–401.

Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24.

Kay, H. (1992). A qualitative model of the space shuttle reaction control system. Technical Report AI92-188, The University of Texas at Austin.

Kraan, I. C., Richards, B. L., and Kuipers, B. J. (1991). Automatic abduction of qualitative models. In *Proceedings of the Fifth International Workshop on Qualitative Reasoning about Physical Systems*, 295–301.

Kuipers, B. (1986). Qualitative simulation. *Artificial Intelligence*, 29:289–338.

Kuipers, B. (1994). *Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge.* MIT Press. In Press.

Nordhausen, B., and Langley, P. (1993). An integrated framework for empirical discovery. *Machine Learning*, 12:17–47.

Richards, B. L., Kraan, I., and Kuipers, B. J. (1992). Automatic abduction of qualitative models. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, 723–728.

Rickel, J. (1992). Automated modeling for answering prediction questions: Exploiting interaction paths. Technical Report AI92-178, University of Texas at Austin.

Rickel, J., and Porter, B. (1992). Automated modeling for answering prediction questions: Exploiting interaction paths. In *Proceedings of the Sixth International Workshop on Qualitative Reasoning about Physical Systems*, 82–95. Heriot-Watt University. Edinburgh, Scotland.

Soderman, U., and Stromberg, J.-E. (1991). Combining qualitative and quantitative knowledge to generate models of physical systems. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, 1158–1163.