

Chapter 13

Applying qualitative modelling to environmental problems

Simon Parsons

*Department of Computer and Information Science, Brooklyn College, City University of New York, 2900 Bedford Avenue, Brooklyn, 11210 New York
(parsons@sci.brooklyn.cuny.edu)*

Abstract

The analysis of many biochemical engineering problems in environmental modelling is based upon the development and solution of sets of differential equations. A complete analytical solution of such a model requires that every numerical constant in this set of equations is precisely known. This chapter describes some techniques, developed in the field of artificial intelligence, which permit the solution of such sets of equations when some constant values are either unknown, or only known imprecisely. The use of these techniques is illustrated with the use of models of anaerobic fermentors.

Keywords: Qualitative modelling, semiquantitative modelling, order of magnitude reasoning

1 Introduction

The basis of the method by which much of modern engineering, especially in areas such as biochemical and environmental engineering, proceeds is by the identification and solution of models based upon sets of differential equations [Bailey and Ollis, 1986]. These equations describe the dynamic behaviour of a system in terms of various key variables. When the equations are solved, the values of those key variables can be used to describe the state of the system at any instant. In theory, the procedure for developing and solving such a model is as follows. First, the differential equations, based on general knowledge of the biochemical process and the reactions which underly them, are written down. These equations are generic descriptions of the processes. These general descriptions are then instantiated for the particular system being studied by identifying the values of constant terms through appropriate experimental measurements. Then the equations are solved, either analytically or by simulation, to give a description of the specific system of interest.

For many biochemical systems, the first stage is often relatively simple. Laws of nature such as the law of mass conservation, are usually sufficient to form the foundation of the necessary set of equations. As a result, given a particular system, it is often relatively easy to collect together a set of differential equations which can form the nucleus of a detailed mathematical model of that system. However, to take this nucleus and flesh it out with all the necessary information that will make it an accurate and realistic working model is far from simple.

This is because the precise values of the constants that relate the variables in the differential equations are often hard to establish. Real environmental systems are horribly complicated and as a result, it is extremely difficult to measure the various constants with any accuracy. This is particularly true when the dynamic behaviour of such systems is considered. Indeed, the systems may be subject to complex relations with their surroundings [Serra *et al.*, 1992] which may make it nearly impossible to isolate them without distorting any measurements made. As a result, it is likely to be extremely difficult, time-consuming and expensive to identify the value of every numerical constant [Steyer *et al.*, 1992]. In practice, therefore, rather than a full set of precise numerical constants, one is faced with an incomplete set of imprecisely known values. Without the values of all the constants the set of equations have no practical value because without known values it is not possible to run a classical simulation, and for most real systems the set of equations cannot be solved by analytical means.

Thus incomplete and uncertain knowledge of the necessary numerical constants would seem to rule out the use of conventional methods in projects such as modelling the scaling up of laboratory fermentors in order to perform tasks like risk evaluation and cost estimation. It need not, however, prevent the use of artificial intelligence techniques. Several authors have applied rule-based methods, for instance [Baldwin *et al.*, 1993; Liong *et al.*, 1991], especially in the area of control of biochemical processes [Serra *et al.*, 1992; Steyer *et al.*, 1991; Watts and Knight, 1991]. A particularly appropriate set of artificial intelligence techniques are those from the area of qualitative modelling [Davis, 1990; Weld and de Kleer, 1990], since many approaches to qualitative modelling are specifically designed to solve sets of differential equations when the constant terms are either unknown, or only known imprecisely.

This chapter discusses such techniques, and shows how they may be used to provide a battery of approaches to handling the kind of models used in environmental engineering. In particular, this chapter focusses on three different approaches under the qualitative modelling banner¹. Section 2 looks at qualitative reasoning, an approach in which it is possible to handle constants whose values are unknown. Section 3 looks at order of magnitude reasoning, which makes it possible to handle information such as “A is much larger than B”—information which it is often possible to obtain in the absence of exact values. Finally Section 4 considers semiquantitative reasoning, an approach which extends qualitative rea-

¹For clarity, in this chapter I use the term “qualitative modelling” to refer to a whole range of techniques, and the term “qualitative reasoning” to refer to one of these. This distinction is not common practice, and the term “qualitative reasoning” is often used to refer to the whole field.

soning when limited amounts of numerical information are available. Each section gives a brief overview of developments in the area and then discusses one particular approach in some detail, with the aim of giving the reader an appreciation of what it is possible to do with it. To help with the latter task, each section also includes a relatively detailed example of the approach applied to an environmental model.

It should be noted that this chapter is not a definitive survey of qualitative modelling techniques—the subject is far too broad. Neither is it a tutorial on how to use qualitative modelling, either in general, or in the context of environmental engineering. Instead it is intended to provide a short guide to the kinds of things that qualitative modelling makes possible, along with pointers to some of the relevant work.

2 Qualitative reasoning

Historically, the first of the qualitative modelling approaches addressed the problem of dealing with numerical values whose magnitude is completely unknown. These are the approaches which will be discussed in this section.

2.1 Overview of qualitative reasoning

The paper that is always cited as being the foundational work in qualitative reasoning is Hayes' Naive Physics Manifesto [1978] in which he urged practitioners of artificial intelligence to "put away childish things by building large scale formalizations" [Hayes, 1985b]. His suggestion was that real progress in the field would come about by attempting to model a large part of human commonsense knowledge about the real world, and his first attempt created an initial theory of liquid behaviour [Hayes, 1985a]. This work was built upon first order logic, the traditional tool of symbolic artificial intelligence. At the same time, and to some extent as a result of Hayes' proposal, work that modelled complex systems in a way that mirrored the kind of approach adopted by engineers was emerging.

There are, broadly speaking, three strands to this work, all having in common the fact that they deal with abstractions of real numbers into positive, negative and zero valued quantities rather than dealing with numbers themselves. The first approach is that of Kuipers [1984] who takes a set of differential equations, abstracts them to just consider their qualitative impact, and then uses them as a set of constraints on the possible values of the state variables. This approach has been implemented as the QSIM software system [Kuipers, 1986], and the complete line of work by Kuipers and his colleagues until 1994 is summarised in [Kuipers, 1994]. The second approach, taken by de Kleer and Brown [1984] and Williams [1984] is to build libraries of components, each of which has a well defined qualitative behaviour described by sets of qualitative differential equations, and connect components together to build a qualitative model. Some of this work is implemented as the

Table 1: Qualitative addition

\oplus	+	0	-
+	+	+	?
0	+	0	-
-	?	-	-

ENVISION software system which takes its name from the process of “envisonment” by which behaviour is inferred from the structure of the system. The final approach not only models components, but also the processes that they may undergo. Work on this approach is primarily due to Forbus [1984], and is closest in spirit to the work on naive physics. In addition, this approach goes further than the others in allowing sets of objects to have group behaviours over and above their individual ones, thus providing a far richer modelling language.

The core of the first two approaches described above is the idea of qualitative differential equations. Rather than attempting to deal with a mass of numerical data, values are only distinguished as positive (+), zero (0), negative (-), or unknown (?). These values are sufficient to identify many of the interesting features of the behaviour of the important variables in a given system. Briefly, this works as follows. Imagine that we have a very simple system which may be described by the equations:

$$\begin{aligned}\frac{dx}{dt} + x &= k_1 \\ \frac{d^2x}{dt^2} + k_2 &= 0\end{aligned}$$

where k_1 and k_2 are positive constants, and x is a substrate concentration. The qualitative abstraction of these equations, in which all numerical values are replaced by +, 0 or - is:

$$\begin{aligned}\frac{dx}{dt} \oplus x &= + \\ \frac{d^2x}{dt^2} \oplus + &= 0\end{aligned}$$

where \oplus is qualitative addition, as described by Table 1. To solve the pair of equations we look for sets of qualitative values that satisfy them, for instance:

$$\begin{aligned}x &= + \\ \frac{dx}{dt} &= + \\ \frac{d^2x}{dt^2} &= -\end{aligned}$$




Figure 1: The qualitative behaviour of x

In other words, x is positive, its first time derivative is positive, but its second time derivative is negative. This set of values tells us that the behaviour of the concentration over time will be to rise to some limiting value as in Figure 1. We may not know what the limit is, but we do know that the concentration will eventually level off, and this less precise information may be sufficient for some purposes. Clearly if we are trying to establish that the substrate concentration has a maximum value then the information we are able to deduce is quite adequate, and in many cases the fact that we can learn something from qualitative reasoning far outweighs the fact that what we learn is not very detailed.

As qualitative reasoning has been applied, one of the important intuitions to emerge is that the process of reasoning about a physical system involves two tasks. One is building a model to describe whatever physical system is being investigated, and the other is simulating the physical system using the model. The idea of compositional modelling [Falkenhainer and Forbus, 1991] revolves around supplying a library of model fragments which may be composed to form a complete model of the situation in hand, and it has been shown to scale-up on such applications as CyclePad [Forbus and Whalley, 1994; Forbus *et al.*, 1999] which captures a significant body of knowledge on thermodynamic processes.

The problem of which model fragments to compose is a significant one, especially if this task is to be automated [Iwasaki and Levy, 1994]. There is also the issue of how to represent the fragments, and how to turn a set of fragments into a unified model which can then be used as the basis for the simulation task. This aspect of the problem is addressed by Farquhar's Qualitative Physics Compiler (QPC). QPC takes a high-level description of components and compiles them down into a form which can be fed into QSIM which then carries out the simulation. QPC has been shown to work well in a wide variety of domains, including socio-economic allocation [Brajnik and Lines, 1998], chemical engineering [Catino, 1993], water supply control [Farquhar and Brajnik, 1994], and plant physiology [Rickel and Porter, 1994].

Having established the kind of problems that can be solved using qualitative reasoning, and the results it is possible to obtain, we will look at one qualitative approach in more detail.

2.2 The QSIM system

The system we will consider is Kuipers' QSIM system [Kuipers, 1986, 1994], which provides a comprehensive suite of programs for evaluating qualitative models. This section gives an overview of the way that models are built and evaluated in QSIM, drawn from the description given in [Cem Say, 1998].

As with any system for qualitative reasoning, the basic component of a QSIM model is a set of state variables. In QSIM these are functions which are continuously differentiable with respect to time. Each variable has a set of values associated with it—this is the *quantity space* for that variable, the set of possible values that the variable can take. These values are also known as qualitative magnitudes, and the usual quantity space is, as already discussed, the set of signs $\{+, 0, -\}$. (The symbol $?$ is not, strictly speaking, part of the quantity space, but an abbreviation for “+ or 0 or -”. Similarly, $(0, +)$ is an abbreviation for “+ or 0”, and $(0, -)$ is an abbreviation for “0 or -”.) We are also interested in the qualitative direction of a variable, that is the sign of its first derivative. These again are typically drawn from a quantity space of $\{+, 0, -\}$, also written as *inc*, *std* and *dec*. The qualitative magnitude and qualitative direction, taken as a pair, make up the *qualitative value* of a variable, and the full set of qualitative values of all the relevant variables defines the state of the system in question.

The second main component of a QSIM model is a set of constraints which capture time-independent relations between variables. Thus, for example, we would represent the fact that:

$$A = BC$$

by using the *mult* constraint:

$$(mult\ A\ B\ C)$$

Broadly speaking there are seven different types of constraint that may be employed in QSIM models. Three capture simple arithmetic relationships—*add*, *minus* and *mult*—and need no further explanation. The constraint *constant* is also fairly simple, specifying the value of a constant term, and d/dt is used to identify variables which are derivatives of other variables. The final two constraints, M^+ and M^- , are a little more complex. They represent the fact that there is a monotonic function relating the two variables named in the constraint. This function is increasing when the constraint is M^+ and decreasing when the function is M^- . The constraints, of course, are just a way of specifying the differential equations which make up the model being encoded in QSIM.

The final component of a QSIM model is the initial state of the system. This is simply a set of initial qualitative values for each variable in the model. If this model is only partially specified, QSIM can generate a set of initial states which are consistent with the partial specifications.

Once the model has been created, QSIM can be run to simulate the behaviour of the system. Starting from the initial state (or set of states), QSIM uses a set of transition rules to

generate all possible new states. Each of these new states is a set of new qualitative values for all the variables. For example, if the qualitative magnitude of a variable is $+$ and its qualitative direction is *dec*, then at some point in the future its magnitude will be 0. Thus the next state has 0 as the qualitative magnitude of that variable. The transitions rules, in general, generate a set of possible next states and each of these generates a further set of next states. The full set of sets of states can be considered as a tree-like structure of trajectories, a set of transitions from one state to another, with some states occurring on several trajectories. The states along a single trajectory make up what is known as a qualitative behaviour.

The role of the constraints in QSIM is to prune the number of states, thus ruling out certain behaviours. Each time the transition rules are applied, the resulting set of states is checked against the constraints. When a state has an assignment of qualitative values to variables which violate one or more constraints, that state is deleted. This ensures that the behaviours generated are consistent with the initial model and so represent legal solutions to the initial set of differential equations which make up the model. Inference in QSIM is thus a “generate and test” process.

QSIM has been widely used by researchers in qualitative reasoning, and, as a result, a number of enhancements have been developed. In general, there are two types of enhancement. One makes it possible to incorporate more detailed information into the models, thus allowing more precise predictions to be obtained². This type thus builds on top of QSIM. The other refines the kind of reasoning carried out in QSIM, eliminating some of the possible behaviours that QSIM generates. This type thus filters the output of the QSIM inference engine. An example of the first kind of enhancement is given by Neitzke and Neumann [Neitzke and Neumann, 1994] in the form of the system RSIM+. RSIM+ allows additional constraints between variables to be stated. For instance, the M^+ constraint is refined to make it possible to identify monotonically increasing constraints which are sub-linear, linear and super-linear. An example of the second kind of system is Cem Say’s qualitative version of L’Hôpital’s filter [Cem Say, 1998]. This checks QSIM-generated behaviours against L’Hôpital’s rule and rules out those which violate it.

2.3 An example of qualitative reasoning

To illustrate the use of qualitative reasoning more completely, consider the following example. It is possible [Bailey and Ollis, 1986] to write down a complex set of equations which fully describe the action of an anaerobic fermentor and which, when solved, provide a suitable model of its behaviour. Unfortunately, the results of this analysis hinge upon the values of a number of key constants whose values not only vary from fermentor to fermentor, but are also extremely difficult to measure. As a result it is difficult and expensive to provide accurate solutions from a conventional analysis. A qualitative analysis is, however,

²As will become apparent, this is a general trend in qualitative reasoning and is the motivation behind the order of magnitude and semi-qualitative techniques discussed later in this chapter.

possible.

The following set of equations provide a simplified model of the behaviour of an anaerobic fermentor:

$$\frac{dx_1}{dt} + (k_{12} + k_{13})x_1 + k_{11}x_1x_5 = k_{21}x_2 \quad (1)$$

$$\frac{dx_2}{dt} + k_{21}x_2 = k_{12}x_1 \quad (2)$$

$$\frac{dx_3}{dt} - k_{63}x_4 = k_{13}x_1 \quad (3)$$

$$\frac{dx_4}{dt} + k_{43}x_4 = k_{11}x_1x_5 \quad (4)$$

$$\frac{dx_5}{dt} + k_{53}x_5 + k_{11}x_1x_5 = 0 \quad (5)$$

where x_1 – x_5 are concentrations of various substrates, either those wastes being digested or the products of the digestion. The k_i are constants. The full model may also be solved using qualitative methods, but the additional detail adds nothing to the understanding of the technique. It is, of course, perfectly possible to apply the method to any equation based model. This particular model was chosen because it was easily available.

To apply QSIM to this model we first need to define the quantity spaces for the variables, writing X1 for x_1 , DX1 for dx_1/dt and so on, we have, in QSIM notation:

(quantity-spaces

- (X1 (minf zero inf) "X1")
- (X2 (minf zero inf) "X2")
- (X3 (minf zero inf) "X3")
- (X4 (minf zero inf) "X4")
- (X5 (minf zero inf) "X5")
- (DX1 (minf zero inf) "DX1")
- (DX2 (minf zero inf) "DX2")
- (DX3 (minf zero inf) "DX3")
- (DX4 (minf zero inf) "DX4")
- (DX5 (minf zero inf) "DX5"))

which indicates that each variable can have the usual +, 0, or – value, because we only distinguish the landmarks ∞ , 0 and $-\infty$. We then specify the constraints between variables which, because the arithmetic constraints are binary, means that we have to introduce four auxilliary variable X6, X7, X8 and X9. We then have:

(constraints

- ((d/dt X1 DX1))
- ((d/dt X2 DX2))


```

((d/dt X2 DX3))
((d/dt X2 DX4))
((d/dt X2 DX5))
((mult X6 DX4 X4))
((add X7 DX1 X1))
((add X8 X6 X5))
((add X9 DX5 X8))
((add X2 X7 X6))
((add X2 DX2 X2))
((add X1 DX3 X4))
((add X6 DX4 X4))
(constant X9 0))

```

The first five of these define DXn to be the first time derivatives of the Xn , the next four define the auxilliary variables, and the following five then capture (1)–(5) respectively. Finally we need to define the initial states of the variables:

```

X1   = ⟨(0, +), (0, +)⟩
X2   = ⟨+, ?⟩
X3   = ⟨(0, +), ?⟩
X4   = ⟨+, ?⟩
X5   = ⟨+, (0, -)⟩
DX1  = ⟨(0, +), ?⟩
DX2  = ⟨?, ?⟩
DX3  = ⟨?, ?⟩
DX4  = ⟨?, ?⟩
DX5  = ⟨(0, -), 0⟩

```

These, for instance define the initial value of x_1 to be either zero or strictly positive, and to have a time derivative which is zero or strictly positive, while x_4 has an initial value which is strictly positive, and a time derivative which is positive, negative, or zero.

Given this model, QSIM would then generate the full set of behaviours which satisfy the initial set of equations. For example, one such behaviour is given in Table 2. As for the initial state, each column in each half of the table gives the prediction for one of the system variables, $x_1 \dots x_5$ and their first derivatives. The value in each cell of the table gives the qualitative value of the corresponding variable, that is its qualitative magnitude and its qualitative direction, thus the second value in each tuple under X1 gives the value of:

$$\frac{dx_1}{dt}$$

which refers to the same quantity as the first value in each tuple under DX1 (as encoded in

Table 2: A QSIM prediction for the anaerobic fermentor model of (1)–(5).

X1	X2	X3	X4	X5	time
$\langle(0, +), (0, +)\rangle$	$\langle+, ?\rangle$	$\langle(0, +), ?\rangle$	$\langle+, -\rangle$	$\langle+, (0, -)\rangle$	t_0
$\langle(0, +), (0, +)\rangle$	$\langle+, ?\rangle$	$\langle(0, +), ?\rangle$	$\langle 0, 0\rangle$	$\langle+, (0, -)\rangle$	t_1
$\langle(0, +), (0, +)\rangle$	$\langle+, ?\rangle$	$\langle(0, +), ?\rangle$	$\langle 0, +\rangle$	$\langle+, (0, -)\rangle$	t_2
$\langle(0, +), (0, +)\rangle$	$\langle+, ?\rangle$	$\langle(0, +), ?\rangle$	$\langle+, +\rangle$	$\langle+, (0, -)\rangle$	t_3

DX1	DX2	DX3	DX4	DX5	time
$\langle(0, +), ?\rangle$	$\langle?, ?\rangle$	$\langle?, ?\rangle$	$\langle-, +\rangle$	$\langle(0, -), 0\rangle$	t_0
$\langle(0, +), ?\rangle$	$\langle?, ?\rangle$	$\langle?, ?\rangle$	$\langle 0, +\rangle$	$\langle(0, -), 0\rangle$	t_1
$\langle(0, +), ?\rangle$	$\langle?, ?\rangle$	$\langle?, ?\rangle$	$\langle+, +\rangle$	$\langle(0, -), 0\rangle$	t_2
$\langle(0, +), ?\rangle$	$\langle?, ?\rangle$	$\langle?, ?\rangle$	$\langle+, +\rangle$	$\langle(0, -), 0\rangle$	t_3

the d/dt constraints). The second value in each tuple under DX1 gives the value of:

$$\frac{d^2x_1}{dt^2}$$

The particular behaviour given in Table 2 traces changes in x_4 while the remaining variables remain in their initial state. This behaviour is a possible, if unlikely, way that the system, as described by the model given above, might evolve. x_4 is initially positive, with a negative first time derivative, and a positive second time derivative (values which are consistent with the initial state). The first time point on the behaviour is when the value of x_4 becomes zero and the first time derivative becomes zero (for the purposes of this example we look at the situation in which these two events are simultaneous). The next time point is the next point at which some value changes, in this case it is the first time derivative of x_4 , which now becomes positive. Finally, under the influence of this positive derivative, x_4 becomes positive again.

QSIM will generate all such behaviours which are consistent with the initial state and the constraints on the variables, following them for as long as desired by the user.

3 Order of magnitude reasoning

Despite the undoubted success of qualitative methods, there are some problems with qualitative reasoning that make it unsuitable for modelling certain systems. These problems stem from the limited number of values that any constant or variable can adopt. In this section we consider order of magnitude techniques, the development of which was motivated by the desire to overcome these limitations of qualitative reasoning.




Figure 2: Two colliding masses

3.1 An overview of order of magnitude reasoning

In the first paper on order of magnitude reasoning, Raiman [1986] illustrated the problems of qualitative reasoning with a simple example from mechanics. Consider two masses which collide while travelling towards one another along the same line (Figure 2). One has a large mass M and velocity V , the other has a small mass and velocity m and v . The net momentum from left to right above is given by the law of the conservation of momentum as:

$$MV_{net} = MV - mv$$

Since M , V , m and v are all positive values, they all have qualitative value $+$, and the net rightwards momentum is established by the calculation:

$$MV_{net} = + \otimes + \ominus + \otimes +$$

where \ominus is the operator representing the difference of two qualitative values (see Table 3), and \otimes is the operator representing the product of two such values (see Table 4). It is clear from Table 4 that the product of two positive values will itself be positive so that the calculation reduces to:

$$MV_{net} = + \ominus +$$

Now, Table 3 summarises the fact that the difference of two values which are only known to be positive can be either positive, negative or zero, depending on the relative sizes of the values. Thus qualitative reasoning can only deduce that the overall rightwards momentum will be $?$, while intuitively we can see that it will be $+$ because MV is much larger than mv .

Now, to some extent this problem is a straw man. Certainly systems such as QSIM provide a way around it. Faced with such a situation, QSIM would identify that the result of the subtraction would depend upon the relative magnitudes of MV and mv . Rather than return the overall momentum as $?$, the system would introduce the quantity $MV - mv$ as a landmark value and give three possible behaviours, one for each qualitative value of the landmark. Thus it would predict that if $MV - mv$ is positive, the net momentum would be rightwards, if $MV - mv$ is negative, the momentum would be leftwards, and if $MV - mv$ is zero, the net momentum would be zero. However, Raiman's main point is still correct—resolving this problem, even with the branching behaviours, involves stepping outside the qualitative reasoning and selecting one output from many (and a typical qualitative model will throw out many such branch points). What would be useful is a way of automating the reasoning, which, after all, is something that humans find quite easy to do.

Table 3: Qualitative subtraction

\ominus	+	0	-
+	?	+	+
0	-	0	+
-	-	-	?

Table 4: Qualitative multiplication

\otimes	+	0	-
+	+	0	-
0	0	0	0
-	-	0	+

The first attempt to provide this automation was proposed by Raiman [1986] in the paper in which he pointed out the problem. He introduced a system called FOG which allowed the representation of “order of magnitude” concepts. Thus it allows the statement that, for instance, A is negligible with respect to B , $A \text{ Ne } B$, or that A has the same sign and order of magnitude as B , $A \text{ Co } B$. These relations are then used to define a set of inference rules such as:

$$A \text{ Ne } B, B \text{ Co } C \Rightarrow A \text{ Ne } C$$

So that if A is much smaller than B , which is about the same size as C , then A is much smaller than C . In all Raiman provides 30 such rules of inference, giving a semantics for the approach which is based on non-standard analysis, and FOG has been used in the modelling of analog circuits [Dague *et al.*, 1987].

FOG has been discussed by Dubois and Prade [1991] who have considered the problem that is caused by the use of non-standard analysis as a basis for a semantics—namely that the results are only valid in the limit. In order to cope with situations in which $A \text{ Co } B$ does not mean that A and B are infinitely close together, they propose a new interpretation in terms of an interval on the ratio of A to B . This allows them to validate the inference rules, and allows a sensible limit on the chaining of inferences such as:

$$\begin{array}{r} 30 \text{ Co } 31 \\ 31 \text{ Co } 32 \\ \hline 30 \text{ Co } 32 \end{array}$$

to be established that prevents the derivation of $30 \text{ Co } 1000$ without the need for an arbitrary

cut-off.

The FOG approach has recently been extended by Dague to provide much more sophisticated approaches to handling order of magnitude reasoning. In particular, he has provided a mechanism for obtaining smooth changes between the orders of magnitude [Dague, 1993b], providing a solution to exactly the same problem addressed by Dubois and Prade. His ROM[K] system, which does this, is based on a set of rules of inference similar to those of FOG, and may be considered a natural successor to it. Dague has also provided a system which can perform the same kind of reasoning using numerical information [Dague, 1993a], and Nayak [Nayak, 1992] has built a rather similar system in which the order of magnitude of a quantity is determined using logarithms. Yip's [Yip, 1996] asymptotic order of magnitude reasoning is also reminiscent of FOG.

Another scheme for order of magnitude reasoning is due to Mavrovouniotis and Stephanopoulos [1987; 1989] who formalised the representation of relations such as $A > B$ to give a system called O[M] that they claim is expressive enough for all engineering problems. The semantics of the relations is provided in terms of the bounds on the ratio between A and B , and two possible interpretations are given. The first is mathematically correct, but conservative, and the second is heuristic but more humanly aggressive in the inferences it sanctions. O[M] has been applied to problems in process engineering [Mavrovouniotis and Stephanopoulos, 1988]. It is possible to show that the O[M] approach can be handled using a formal system for reasoning using intervals [Parsons, 1993], and Travé-Massuyès and Piera have shown that such interval systems can be used as the basis of a general approach to order of magnitude reasoning [Travé-Massuyés and Piera, 1989].

3.2 The O[M] system

In this section we describe O[M], introduced by Mavrovouniotis and Stephanopoulos [1987; 1989] in detail, prior to showing how it may be used to simplify a biochemical model. O[M] is based upon a set of seven primitive relations:

$A \ll B$	A is much smaller than B
$A -< B$	A is moderately smaller than B
$A \sim< B$	A is slightly smaller than B
$A == B$	A is exactly equal to B
$A >\sim B$	A is slightly larger than B
$A >-B$	A is moderately larger than B
$A \gg B$	A is much larger than B

These primitive relations can then be composed to form further relations. The compositions which make sense are those of relations which are consecutive in the above list, so for instance, we have “ A is less than B ”, which is the composition of \ll , $-<$ and $\sim<$. This is denoted by:

$$A \ll \dots \sim< B$$

There are a total of 21 sensible compound relations, and Mavrovouniotis and Stephanopoulos claim that these are sufficient to capture all order of magnitude relations commonly used by engineers. The meaning of these relations are given in terms of the ratio of the two quantities in question. Thus we take $A \ll B$ to mean that:

$$\frac{A}{B} \ll 1$$

and similarly for the other relations. We can then [Parsons, 1993] define how to translate between numerical values of A and B and the O[M] relations by applying the following mappings:

$$[[A, B]]_h = \begin{cases} A \ll B & \text{if } e_1 > A/B \\ A - < B & \text{if } e_5 < A/B < e_2 \\ A \sim < B & \text{if } e_6 < A/B < 1 \\ A == B & \text{if } 1 = A/B \\ A > \sim B & \text{if } 1 < A/B < e_7 \\ A > - B & \text{if } e_3 < A/B < e_8 \\ A \gg B & \text{if } e_4 < A/B \end{cases}$$

which generates a relation from A and B , and:

$$[[A \text{ rel } B]]'_h = \begin{cases} e_5 > A/B & \text{if rel is } \ll \\ e_1 < A/B < e_6 & \text{if rel is } - < \\ e_2 < A/B < 1 & \text{if rel is } \sim < \\ 1 = A/B & \text{if rel is } == \\ 1 < A/B < e_3 & \text{if rel is } > \sim \\ e_7 < A/B < e_4 & \text{if rel is } > - \\ e_8 < A/B & \text{if rel is } \gg \end{cases}$$

which gives the bounds on the relative magnitudes of A and B from a relation. The e_i are parameters chosen by the user and can be described in terms of a single parameter e :

$$\begin{aligned} e_1 &= e \\ e_2 &= \frac{1}{1+e} \\ e_3 &= 1+e \\ e_4 &= \frac{1}{e} \\ e_5 &= \frac{e}{1+e} \\ e_6 &= \frac{1}{(1+e)^2} \\ e_7 &= (1+e)^2 \\ e_8 &= \frac{1+e}{e} \end{aligned}$$

It is entirely intentional that the intervals defined by these mappings either overlap, or are non-contiguous. The idea is that A/B is mapped into a set of overlapping intervals,

is transformed somehow, and then “shrunk back” to the non-contiguous intervals. The reason is to model the fact that people reason more “aggressively” than is warranted by the available information. This works as follows. Consider we are told that:

$$\begin{aligned} A &>\sim B \\ B &>\sim C \end{aligned}$$

then from the second mapping we know that:

$$\begin{aligned} 1 &< A/B < (1 + e) \\ 1 &< B/C < (1 + e) \end{aligned}$$

From this we can infer that:

$$1 < A/C < (1 + e)^2$$

which the first mapping tells us means that:

$$A >\sim C$$

Now, the important thing about the heuristic mapping is that if we then want to use the fact that:

$$C >\sim D$$

to infer something about the relative orders of magnitude of A and D , we use the first mapping again to get:

$$\begin{aligned} 1 &< A/C < (1 + e) \\ 1 &< C/D < (1 + e) \end{aligned}$$

and go through the same procedure as before to get:

$$A >\sim D$$

In the next section we give further examples of the kind of reasoning supported by O[M].

3.3 An example of order of magnitude reasoning

To illustrate one possible use of order of magnitude reasoning, we will use O[M] to simplify another anaerobic fermentor model, this time taken from [Bailey and Ollis, 1986]. The model is given in (6)–(22):

$$\frac{dp_{CO_2}}{dt} = -p_T \frac{V}{\rho_g V_G} T_G - \frac{p_{CO_2}}{V_G} Q \quad (6)$$

$$Q_{CO_2} = \frac{V}{\rho_g} T_G \quad (7)$$

$$Q = Q_{CO_2} + Q_{CH_4} + Q_{H_2O} \quad (8)$$

$$(hs) = \frac{s(h^+)}{K_a} \quad (9)$$

$$(h^+) = \frac{K_1 [CO_2]_D}{[HCO_3^-]} \quad (10)$$

$$[HCO_3^-] = z - s \quad (11)$$

$$\frac{dz}{dt} = \frac{F}{V}(z_0 - z) \quad (12)$$

$$T_G = k_L a ([CO_2]_D^* - [CO_2]_D) \quad (13)$$

$$[CO_2]_D^* = K_H a p_{CO_2} \quad (14)$$

$$\frac{d[tox]}{dt} = \frac{F}{V} ([tox]_0 - [tox]) \quad (15)$$

$$\frac{d[CO_2]_D}{dt} = \frac{F}{V} ([CO_2]_{D0} - [CO_2]_D) + T_G + R_B + R_C \quad (16)$$

$$R_C = \frac{F}{V} ([HCO_3^-]_0 - [HCO_3^-]) + \frac{ds}{dt} + \frac{dz}{dt} \quad (17)$$

$$\frac{dx}{dt} = \frac{F}{V} (x_0 - x) + \mu x - k_T [tox] \quad (18)$$

$$\frac{ds}{dt} = \frac{F}{V} (s_0 - s) - \frac{\mu}{Y_{x/s}} x \quad (19)$$

$$\mu = \frac{\mu_{max}}{\left[1 + \frac{K_s}{(hs)} + \frac{(hs)}{K_i}\right]} \quad (20)$$

$$R_B = Y_{CO_2/x} \mu x \quad (21)$$

$$Q_{CH_4} = \frac{V}{\rho_g} Y_{CH_4/x} \mu x \quad (22)$$

Clearly this is a more complex model than that presented above, though it has the same general form. Because of this complexity, it is rather hard to extract useful information directly from the equations. However, if the equations are simplified, such information can be extracted, and one means of carrying out this simplification is to use order of magnitude reasoning.

Consider (20). One thing we might want to extract from this equation is the relationship between μ and its maximum value μ_{max} in steady state operation. In the steady state, we know³ that (hs) is roughly equal to K_s , but much smaller than K_i . In other words;

$$\begin{aligned} (hs) &\sim \langle \dots \rangle \gg K_s \\ (hs) &\ll K_i \end{aligned}$$

Now, given this initial information, O[M] can be applied to (20) in the same way as we saw above (for the detailed calculation see [Parsons, 1993]) to determine that:

$$\mu_{max} \gg \mu$$

³Note that these examples are constructed with the intention of being instructive in the use of O[M] rather than realistic.

which tells us that the rate is moderately smaller than the maximum.

We can also use O[M] to establish something about the rate of change of $[CO_2]_D$, which is the concentration of dissolved CO_2 . The situation we want to investigate is when the concentration is well above the initial concentration and slightly less than the equilibrium concentration:

$$\begin{aligned} [CO_2]_D &\ll [CO_2]_{D_0} \\ [CO_2]_D &\sim < [CO_2]_D^* \end{aligned}$$

Taking these and combining the first with (16) and the second with (13) gives us, respectively:

$$\begin{aligned} \frac{d[CO_2]_D}{dt} &> \sim \frac{F}{V}[CO_2]_D + (T_G + R_B + R_C) \\ T_G &- < -k_L a [CO_2]_D \end{aligned}$$

Combining these, and adding in the additional knowledge that:

$$k_L a \gg \frac{F}{V}$$

gives us:

$$\frac{d[CO_2]_D}{dt} \ll \dots - < [CO_2]_D k_L a + R_B + R_C$$

which is considerably simpler than the original. Further simplification is possible [Parsons, 1993].

The results of this kind of simplification can be used directly, or as input to a system like QSIM which can solve the simplified model qualitatively. In such a case order of magnitude reasoning can be considered a form of pre-processing which provides a rigorous approach to abstracting away unnecessary detail. Alternatively, this kind of reasoning could be used as a further global filter on the output of QSIM, purging branching behaviours which do not fit in with the result of the order of magnitude analysis. It should be noted that exactly the same kind of reasoning can be performed using Nayak's NAPIER system [Nayak, 1992] or Williams and Raiman's caricatural reasoning [Williams and Raiman, 1994].

4 Semiquantitative reasoning

Order of magnitude reasoning provides one means by which limited amounts of information about the magnitude of constant terms can be taken into account in qualitative models. There are other approaches, all of which integrate, to some extent qualitative and quantitative information including [Dormoy, 1988; Féray Beaumont, 1991; Kuipers and Berleant,




Figure 3: An initial set of semiquantitative intervals

1988; Steyer *et al.*, 1992; Sticklen *et al.*, 1991]. I refer to these as “semiquantitative” approaches⁴. One influential approach is embodied in Berleant and Kuipers’ [Berleant and Kuipers, 1997] Q2 and Q3 which operate as filters on top of QSIM, ensuring that the behaviours generated are consistent with whatever numerical information is available. The same approach was used to create the Semi-Quantitative Physics Compiler [Brajnik, 1994; Farquhar and Brajnik, 1994]. These methods have been used, for example, for carrying out comparative analyses [Vatcheva and de Jong, 1999]. Another interesting approach is to bring in fuzzy information [Bellazzi *et al.*, 1999; Shen and Leitch, 1993; Vescovi and Travé-Massuyés, 1992]. In the remainder of this chapter we present a third approach—a generalisation of qualitative reasoning, known as semiquantitative reasoning.

4.1 An overview of semiquantitative reasoning

In semiquantitative reasoning the values of variables and constants are restricted to a set of $2k + 1$ intervals [Parsons and Dohnal, 1993]. This set of intervals covers all numbers from ∞ to $-\infty$, and the intervals are continuous and non-overlapping, so that any real number falls into one, and only one, interval. The intervals are symmetric about zero, which is a distinguished value, and there are k positive and k negative intervals. The boundaries of the intervals may be set by an arithmetic or geometric progression, or may be chosen to reflect what are considered to be interesting values. Since the set of values used in qualitative reasoning corresponds to the set of semiquantitative intervals obtained for $k = 1$, it is clear that semiquantitative reasoning is a generalisation of qualitative reasoning.

A basic understanding of how semiquantitative reasoning may be used to solve sets of differential equations may be obtained from a simple example. Consider the following set of equations:

$$\begin{aligned}x_1 + x_2 &= x_3 \\x_1 \cdot x_4 &= x_3 \\ \frac{dx_4}{dt} &= x_5\end{aligned}$$

The model is solved just as the qualitative one was by finding a set of values for the variables and their derivatives from the set of all possible values⁵ so that the equations are satisfied. If we have the set of intervals depicted in Figure 3 then the following five triplets describe

⁴Kuipers prefers the term “semiquantitative”.

⁵The set of all possible values is not restricted to the set of $2k + 1$ intervals. All compositions of contiguous intervals are also permitted values.

one set of assignments of values to the five variables, and thus one conceivable state of the system:

	x	$\frac{dx}{dt}$	$\frac{d^2x}{dt^2}$
x_1	$[0, 20]$	$[0, 10]$	0
x_2	$[20, 100]$	$[0, 10]$	$[-20, -100]$
x_3	$[10, 20]$	$[20, 100]$	$[20, 100]$
x_4	$[0, 10]$	$[0, 10]$	0
x_5	$[10, 20]$	$[0, -10]$	0

This state is not however a physically possible state of the system since it is not a solution of equations that describe the system. This is because x_3 is determined from $x_1 = [0, 20]$ and $x_2 = [20, 100]$ from the first equation which gives a value of:

$$[0, 20] \oplus [20, 100] \mapsto [20, 500]$$

where \oplus gives the result of adding two intervals using interval arithmetic [Moore, 1966] (in this case $[20, 120]$) and then finding the smallest interval or composition of intervals that holds the result. This value of x_3 contrasts with that of the proposed solution, in which $x_3 = [10, 20]$, and this contradiction rules out the solution. By similar means, a generate and test approach can identify all the sets of five triplets which are solutions of the set of equations, and these correspond to all the semiquantitative states of the model.

By allowing variables to take on a wider range of values than qualitative reasoning does, semiquantitative reasoning permits the use of those numerical values that are known, and this means that it generates more precise solutions than are possible using qualitative reasoning. However, the fact that it is not necessary to have any more information than whether a quantity is positive, negative or zero means that, as is the case for qualitative reasoning, semiquantitative reasoning is very robust, and may be used in situations where conventional methods cannot be used.

4.2 The QSENECA system

Having seen the kind of results that semiquantitative analysis can generate, we consider a software system which can perform a semiquantitative analysis on a set of differential equations. Since it is described at length in [Parsons and Dohnal, 1995] we just give a brief description here. Overall, the program works in a very similar way to QSIM. It takes the relations between the semiquantitative variables of the set of equations to be a series of constraints upon their value. The analysis then consists of taking the known values, and propagating these through the network of constraints, seeing how they affect those values that are initially undefined.

Input to the program is in four parts. The first is the set of differential equations describing the model to be analysed. As for QSIM, these equations are expressed in terms of functions

Table 5: Functional blocks, note that DX is short for $\frac{dx}{dt}$, and DDX for $\frac{d^2x}{dt^2}$

Identification	Description
<i>M1</i>	Addition
<i>M2</i>	Multiplication
<i>M3</i>	Derivative
<i>M4</i>	$X > Y$
<i>M5</i>	$DX > DY$
<i>M6</i>	$DDX > DDY$

relating the variables of interest—here the functions are described as functional blocks such as those in Table 5, and the equations are considered as a graph of connected functional blocks. The second part of the input is the set of semiquantitative intervals to be used for the analysis. The third part of the input is the list of variables whose value is required in the output. All of these parts of the input are fairly straightforward. The final part of the input is more complex. One way of looking at it is as a form of query about the system, another is as a set of constraints on the output. Either way, it is a list of variables and the intervals into which their values fall, and, as just mentioned, we can take this to be a way of ruling out solutions which do not agree with these values, or as a request to find states of the model which agree with them.

Once given this information the program first compiles a set of combinator tables from the set of semiquantitative intervals—one of the overheads of the semiquantitative approach is the need to build a combinator table for each functional block every time that the program is run. Next the program decides on an order in which to test the values of the variables. This is done in such a way that the most constrained variable has its value propagated first, so that once its value is established the conceivable values of all the related variables may be evaluated as swiftly as possible. After these two steps the program begins the process of propagating the constraints, essentially following a generate and test strategy for the value and derivatives of every variable. Finally, after applying all the constraints to all the variables and establishing their possible values, the system outputs a list of all the possible interval values of all the derivatives of all the variables listed in the final part of the input.

4.3 An example of semiquantitative reasoning

As an example of the use of the QSENECA system, consider the semiquantitative analysis of the same anaerobic fermentor as we used in Section 2, repeated here for convenience:

$$\frac{dx_1}{dt} + (k_{12} + k_{13})x_1 + k_{11}x_1x_5 = k_{21}x_2$$

$$\begin{aligned}\frac{dx_2}{dt} + k_{21}x_2 &= k_{12}x_1 \\ \frac{dx_3}{dt} - k_{63}x_4 &= k_{13}x_1 \\ \frac{dx_4}{dt} + k_{43}x_4 &= k_{11}x_1x_5 \\ \frac{dx_5}{dt} + k_{53}x_5 + k_{11}x_1x_5 &= 0\end{aligned}$$

As with the qualitative model, the fact that a specific example is used should not distract attention from the fact that the method can be used to solve any equation based model. The model used for the semiquantitative analysis consists of the same set of differential equations as before, plus those numerical values that are known. These are the values of the following constants:

$$\begin{aligned}k_{11} &= 100 \\ k_{12} &= 1.5 \\ k_{13} &= 5.0 \\ k_{21} &= 3.0 \\ k_{43} &= 1.0 \\ k_{53} &= 0.3 \\ k_{63} &= -1.0\end{aligned}$$

The first stage in the analysis is to write the equations that describe the model in form in which they may easily be specified using the functional blocks. Initially they are written as a series of variables related only by addition and equality. There is no subtraction block since subtraction causes problems in interval arithmetic, and any equation written using subtraction may be rewritten using addition. This generates a new set of equations:

$$\begin{aligned}x_6 + x_{11} + x_{19} &= x_{12} \\ x_7 + x_{12} &= x_{14} \\ x_{15} + x_{16} &= x_8 \\ x_9 + x_{16} &= x_{19} \\ x_{10} + x_{17} + x_{19} &= x_{18}\end{aligned}$$

This second set of equations may be directly written down in terms of functional blocks. There are further equations which relate the variables in the above equations to each other, just as in the QSIM model:

$$\begin{aligned}x_6 &= \frac{dx_1}{dt} \\ x_7 &= \frac{dx_2}{dt} \\ x_8 &= \frac{dx_3}{dt}\end{aligned}$$

$$\begin{aligned}
x_9 &= \frac{dx_4}{dt} \\
x_{10} &= \frac{dx_5}{dt} \\
x_{11} &= 6.5x_1 \\
x_{12} &= 3.0x_2 \\
x_{13} &= 100x_1 \\
x_{14} &= 1.5x_1 \\
x_{15} &= 5.0x_1 \\
x_{16} &= 1.0x_4 \\
x_{17} &= 0.3x_5 \\
x_{18} &= 0 \\
x_{19} &= x_{13}x_5 \\
x_{20} &= x_{19} + x_{11} \\
x_{21} &= x_{19} + x_{17}
\end{aligned}$$

Having done this, it is simple to transform the set of equations to a network of functional blocks, and this comprises the first part of the input to the program. The second part of the input are the semiquantitative intervals, which are those of Figure 4. These are a default set of boundaries suitable for a first attempt at an analysis. The third part of the input is the list of variables whose value are required in the output. In the example, since we are interested in the values of x_2 – x_4 this part of the input will contain the names of these variables and their first and second derivatives, because we want to know the value of all three. The fourth part of the input is a set of additional constraints given by stating the values of some of the variables and their derivatives:

	x	$\frac{dx}{dt}$	$\frac{d^2x}{dt^2}$
x_1	[0, 20]	[0, 10]	0
x_2	[10, 1000]	?	?
x_3	[0, 20]	?	?
x_4	[10, 1000]	?	?
x_5	[10, 20]	[0, -10]	0

Note that ? is shorthand for the interval $[-\infty, \infty]$. This set of constraints may be considered as a query, in this case asking the question:

When x_1 is present in a concentration of less than 20, what are the ways in which it is possible to achieve a linear ($\frac{d^2x_1}{dt^2} = 0$) increase of concentration of x_1 of less than 10 units per unit time while x_5 is present with a concentration of between 10 and 20, and changes linearly ($\frac{d^2x_5}{dt^2} = 0$) at a rate of less than 10




Figure 4: A second set of semiqualitative intervals

units per unit time? Meanwhile x_3 is known to have a positive concentration of less than 20, while that of x_2 and x_4 is between 10 and 1000. The way that these last three variables change with time is not known.

This is, of course, a more constrained version of the problem solved by QSIM in Section 2. Solving this model gives the following as one possible behaviour (we only write down the qualitative state of x_2 , x_3 and x_4 , the variables whose time derivatives are unknown):

	x	$\frac{dx}{dt}$	$\frac{d^2x}{dt^2}$
x_2	[20, 100]	[20, 100]	[0, -10]
x_3	[10, 20]	[500, 1000]	[-20, -100]
x_4	[500, 1000]	[-500, -1000]	[500, 1000]

which gives us a reasonably detailed idea of what values the substrate concentrations might have, and less detailed but still useful information on how they might change over time. Of course, in general, there will be a number of semiqualitative solutions for every qualitative one, each with slightly different interval values for variables and their derivatives. These are, of course, the result of various branching histories in the semiqualitative analysis.

The analysis may be refined. For instance, if we want to further investigate the value and first derivative of x_2 , say, we could choose a new set of intervals, choosing the upper limit of the third positive interval to be 50 instead of 100 as in Figure 5. With the rest of the input the same as before, the following behaviour is one generated:

	x	$\frac{dx}{dt}$	$\frac{d^2x}{dt^2}$
x_2	[20, 50]	[20, 50]	[0, -10]
x_3	[10, 20]	[500, 1000]	[-20, -50]
x_4	[500, 1000]	[-500, -1000]	[500, 1000]

which shows that by making certain intervals narrower, it is possible to make the solution more accurate, in that the intervals in the solution become narrower also. This process could, of course, be repeated. We could split the interval [20, 50] into [20, 35] and [35, 50] in order to further narrow down the possible values of x_2 , or split [500, 1000] into [500, 750] and [750, 1000] to get a better idea of the possible values of x_4 . Of course, every such division increases the computation needed and may well generate a large number of additional solutions.




Figure 5: A third set of semiqualitative intervals

5 Summary

The modelling of environmental problems often involves the solution of models, expressed in terms of differential equations, which have missing or partially known numerical values. This makes it difficult to solve the models using conventional simulation techniques. This chapter has discussed the use of a range of techniques from artificial intelligence which provide ways of solving such models. While the techniques do not overcome the difficulties entirely—since it is not possible, in general, to provide precise answers from imprecise and incomplete data—they do provide means of obtaining some solutions which may be useful.

The chapter has introduced three strands of work from the area of qualitative modelling which seem particularly applicable for environmental engineering—qualitative reasoning, order of magnitude reasoning, and semiqualitative reasoning. For each technique the chapter has given an overview of the style of reasoning it provides, a description of a particular system which provides that kind of reasoning, and an illustration of that kind of reasoning applied to a model of an anaerobic fermentor. Despite this concentration on one type of model, it is hopefully clear that the techniques described could be applied to a wide range of environmental problems.

References

- J. E. Bailey and D. F. Ollis. *Biochemical Engineering Fundamentals*. McGraw Hill, New York, 1986.
- J. F. Baldwin, T. P. Martin, and Y. Zhou. A FRIL knowledge base for the management of uncertainty in performance assessment of hazardous waste repositories. ITRC 188, Advanced Computing Research Centre, University of Bristol, 1993.
- R. Bellazzi, R. Guglielmann, and L. Ironi. A qualitative-fuzzy framework for nonlinear black-box system identification. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, 1999.
- D. Berleant and B. Kuipers. Qualitative and quantitative simulation: bridging the gap. *Artificial Intelligence*, 95:215–256, 1997.
- G. Brajnik and M. Lines. Qualitative modeling and simulation of socio-economic phenomena. *Journal of Artificial Societies and Social Simulation*, 1, 1998.
- G. Brajnik. Semi-quantitative physics compiler: where are we now and where are we going. Technical Report UDMI-RT-03-94, Dipartimento di Matematica e Informatica, University of Udine, 1994.
- C. Catino. *Automated modeling of chemical plants with application to hazard and operability studies*. PhD thesis, Department of Chemical Engineering, University of Pennsylvania, 1993.
- A. C. Cem Say. L'Hôpital's filter for QSIM. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20:1–8, 1998.
- P. Dague, O. Raiman, and P. Devès. Troubleshooting: when modelling is the trouble. In *Proceedings of the 6th National Conference on Artificial Intelligence*, Seattle, WA., 1987.
- P. Dague. Numeric reasoning with relative orders of magnitude. In *Proceedings of the 11th National Conference on Artificial Intelligence*, Washington D. C., 1993.
- P. Dague. Symbolic reasoning with relative orders of magnitude. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, Chambéry, France, 1993.
- E. Davis. *Representations of commonsense knowledge*. Morgan Kaufmann, San Mateo, CA., 1990.
- J. de Kleer and J. S. Brown. A qualitative physics based on confluences. *Artificial Intelligence*, 24:7–83, 1984.
- J-L. Dormoy. Assembling a device. *Artificial Intelligence in Engineering*, 3:216–226, 1988.
- D. Dubois and H. Prade. Semantic considerations on order of magnitude reasoning. In M. G. Singh and L. Travé-Massuyés, editors, *Decision Support Systems and Qualitative Reasoning*. Elsevier Science Publishers, Amsterdam, 1991.
- B. Falkenhainer and K. D. Forbus. Compositional modeling: Finding the right model for the job. *Artificial Intelligence*, 51:95–143, 1991.
- A. Farquhar and G. Brajnik. A semi-quantitative physics complier. In *Proceedings of the 8th International Workshop on Qualitative Reasoning*, 1994.

- S. Féray Beaumont. Qualitative model-based reasoning: an application to distillation process. In M. G. Singh and L. Travé-Massuyés, editors, *Decision Support Systems and Qualitative Reasoning*. Elsevier Science Publishers, Amsterdam, 1991.
- K. D. Forbus and P. B. Whalley. Using qualitative physics to build articulate software for thermodynamics education. In *Proceedings of the 12th National Conference on Artificial Intelligence*, Menlo Park CA/Cambridge MA, 1994. AAAI Press/MIT Press.
- K. D. Forbus, P. B. Whalley, J. O. Everett, L. Ureel, M. Brokowski, J. Baher, and S. E. Kuehne. CyclePad: an articulate virtual laboratory for engineering thermodynamics. *Artificial Intelligence*, 114:297–347, 1999.
- K. D. Forbus. Qualitative process theory. *Artificial Intelligence*, 24:85–168, 1984.
- P. Hayes. The naive physics manifesto. In D. Michie, editor, *Expert Systems in the Microelectronic Age*. Edinburgh University Press, Edinburgh, 1978.
- P. Hayes. Naive physics I: ontology for liquids. In J. R. Hobbs and R. C. Moore, editors, *Formal Theories of the Commonsense World*. Ablex Publishing Corporation, Norwood, NJ., 1985.
- P. Hayes. The second naive physics manifesto. In J. R. Hobbs and R. C. Moore, editors, *Formal Theories of the Commonsense World*. Ablex Publishing Corporation, Norwood, NJ., 1985.
- Y. Iwasaki and A. Y. Levy. Automated model selection for simulation. In *Proceedings of the 12th National Conference on Artificial Intelligence*, Menlo Park CA/Cambridge MA, 1994. AAAI Press/MIT Press.
- B. Kuipers and D. Berleant. Using incomplete quantitative knowledge in qualitative reasoning. In *Proceedings of the 7th National Conference on Artificial Intelligence*, Saint Paul, MN, 1988.
- B. Kuipers. Commonsense reasoning about causality: deriving behaviour from structure. *Artificial Intelligence*, 24:169–203, 1984.
- B. Kuipers. Qualitative simulation. *Artificial Intelligence*, 29:289–338, 1986.
- B. Kuipers. *Qualitative reasoning*. MIT Press, Cambridge, MA., 1994.
- S. Y. Liong, W. T. Chou, and C. C. Law. An expert system for storm water management modelling and its application. *Engineering Applications of Artificial Intelligence*, 3:367–375, 1991.
- M. L. Mavrovouniotis and G. Stephanopoulos. Reasoning with orders of magnitude and approximate relations. In *Proceedings of the 6th National Conference on Artificial Intelligence*, Seattle, WA., 1987.
- M. L. Mavrovouniotis and G. Stephanopoulos. Formal order-of-magnitude reasoning in process engineering. *Computers in Chemical Engineering*, 12:867–880, 1988.
- M. L. Mavrovouniotis and G. Stephanopoulos. Order-of-magnitude reasoning with O[M]. *Artificial Intelligence in Engineering*, 4:106–114, 1989.
- R. E. Moore. *Interval Analysis*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1966.
- P. P. Nayak. Order of magnitude reasoning using logarithms. In *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, Cambridge, MA, 1992.
- M. Neitzke and B. Neumann. Comparative simulation. In *Proceedings of the 12th National Conference on Artificial Intelligence*, Menlo Park CA/Cambridge MA, 1994. AAAI Press/MIT Press.

- S. Parsons and M. Dohnal. Qualitative, semiquantitative and interval algebras and their application to engineering problems. *Engineering Applications of Artificial Intelligence*, 5:553–560, 1993.
- S. Parsons and M. Dohnal. The qualitative and semi-qualitative analysis of environmental problems. *Environmental Software*, 10:75–85, 1995.
- S. Parsons. Using interval algebras to model order of magnitude reasoning. *Artificial Intelligence in Engineering*, 8:87–98, 1993.
- O. Raiman. Order of magnitude reasoning. In *Proceedings of the 5th National Conference on Artificial Intelligence*, Philadelphia, PA, 1986.
- J. Rickel and B. Porter. Automated modeling for answering prediction questions: selecting the time scale and system boundary. In *Proceedings of the 12th National Conference on Artificial Intelligence*, Menlo Park CA/Cambridge MA, 1994. AAAI Press/MIT Press.
- P. Serra, J. Lafuenta, R. Mareno, C. de Prada, and M. Poch. Development of a realtime expert system for wastewater treatment plants. In *Proceedings of the IFAC/IFIP/IMACS International Symposium on AI in Real Time Control*, Delft, The Netherlands, 1992.
- Q. Shen and R. R. Leitch. Fuzzy qualitative simulation. *IEEE Transactions on Systems, Man and Cybernetics*, 23:1038–1061, 1993.
- J. P. Steyer, J. B. Pourciel, D. Simoes, and J. L. Uribelarrea. Qualitative knowledge modelling used in a real time expert system for biotechnological process control. In M. G. Singh and L. Travé-Massuyés, editors, *Decision Support Systems and Qualitative Reasoning*. Elsevier Science Publishers, Amsterdam, 1991.
- J. P. Steyer, I. Queinnec, and D. Simoes. Biotech: a real time application of artificial intelligence for fermentation processes. In *Proceedings of the IFAC/IFIP/IMACS International Symposium on AI in Real Time Control*, Delft, The Netherlands, 1992.
- J. Sticklen, A. Kamel, and W. E. Bond. Integrating quantitative and qualitative computations in a functional framework. *Engineering Applications of Artificial Intelligence*, 4:1–10, 1991.
- L. Travé-Massuyés and N. Piera. The orders of magnitude models as qualitative algebras. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, MI., 1989.
- I. Vatcheva and H. de Jong. Semi-quantitative comparative analysis. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, 1999.
- M. R. Vescovi and L. Travé-Massuyés. A constructive approach to qualitative fuzzy simulation. In *Proceedings of the 6th International Workshop on Qualitative Reasoning*, 1992.
- P. Watts and B. Knight. Fault diagnosis in ASPEX: an expert system for the control of the activated sludge process. *Engineering Applications of Artificial Intelligence*, 4:151–155, 1991.
- D. Weld and J. de Kleer. *Readings in qualitative reasoning about physical systems*. Morgan Kaufmann, San Mateo, CA., 1990.
- B. C. Williams and O. Raiman. Decompositional modeling through caricatural reasoning. In *Proceedings of the 12th National Conference on Artificial Intelligence*, Menlo Park CA/Cambridge MA, 1994. AAAI Press/MIT Press.
- B. C. Williams. Qualitative analysis of MOS circuits. *Artificial Intelligence*, 24:281–346, 1984.

K. M. Yip. Model simplification by asymptotic order of magnitude reasoning. *Artificial Intelligence*, 80:309–348, 1996.