



MONET2

Project Full Title: Network of Excellence on Model-based Systems and Qualitative Reasoning.

Contract: Concerted Action / Thematic Network

Contract No: IST-33540

Deliverable B5:

A Report on Synergies with Model-based Technology and other AI Approaches in Biomedicine

Date: 10 November 2004

Version: 3.0

Status: [Release]

MONET Project Office
Department of Computer Science
University of Wales
Aberystwyth
Ceredigion
SY23 3DB

Contents

| | |
|---|-----------|
| CONTENTS | 2 |
| AUTHORS | 3 |
| 1 EXECUTIVE SUMMARY | 4 |
| 2 INTRODUCTION..... | 4 |
| 2.1 EDITORIAL HISTORY | 5 |
| 3 CAUSAL REASONING AND THEIR QUANTITATIVE ANALOGUES | 5 |
| 3.1 CAUSAL KNOWLEDGE IN BIOMEDICINE | 6 |
| 3.2 ABDUCTIVE REASONING AND DIAGNOSIS | 6 |
| 3.3 BAYESIAN NETWORKS FOR MODEL-BASED DIAGNOSIS..... | 8 |
| 3.3.1 Qualitative probabilistic networks | 8 |
| 3.3.2 Bayesian networks and maximum a priori assignment diagnosis | 10 |
| 3.4 THE RELATIONSHIP BETWEEN QUALITATIVE AND BAYESIAN APPROACHES | 12 |
| 4 QUALITATIVE REASONING AND FUZZY LOGIC..... | 12 |
| 4.1 BACKGROUND | 13 |
| 4.2 SYNERGIES BETWEEN QUALITATIVE MODELS AND FUZZY SYSTEMS | 14 |
| 4.2.1 Fuzzy Qualitative Simulation | 14 |
| 5 QUALITATIVE REASONING AND INDUCTIVE LOGIC PROGRAMMING..... | 18 |
| 6 SYNERGIES BETWEEN MATHEMATICAL MODELS AND QUALITATIVE METHODOLOGIES | 21 |
| 6.1 QR-BASED AUTOMATED MODELLING: NEW PERSPECTIVES IN PHARMACEUTICAL DESIGN | 22 |
| 6.2 QR-BASED AUTOMATED DATA INTERPRETATION: NEW PERSPECTIVES IN ELECTROCARDIOGRAPHIC IMAGING | 23 |
| 7 FUTURE DIRECTIONS IN SYNERGY | 24 |
| 8 REFERENCES..... | 24 |
| 9 DOCUMENT HISTORY | 28 |

Authors

| | |
|----------------|-----------------------------|
| George Coghill | University of Aberdeen |
| Peter Lucas | Radboud University Nijmegen |
| Liliana Ironi | IMATI - CNR |
| René Quiniou | IRISA |
| Iain Russell | UWA, MONET Project |

1 Executive Summary

Symbolic representation and reasoning approaches underlie many systems in artificial intelligence (AI), in particular this holds for model-based systems using qualitative reasoning technology. Yet, the principles underlying qualitative reasoning and model-based systems have their analogues in areas of artificial intelligence where there is an emphasis on quantitative approaches. However, the distinction between the two approaches, i.e. quantitative and qualitative, is sometimes not sharp, as there are methods underlying model-based systems that profit from combining qualitative and quantitative methods. This implies that synergies between model-based systems and qualitative reasoning and other, in particular quantitative methods, in AI exist. However, in this report we also consider related AI methods, such as inductive logic programming.

In the report, we discuss such synergies, focusing on causal knowledge and its use through abductive reasoning, on qualitative reasoning and fuzzy logic, on learning using inductive logic programming and finally on the development of mathematical models using qualitative methodologies. It is believed that the existence of such synergies offer good opportunities for model-based systems and qualitative reasoning technology to be taken up by fields outside the traditional model-based systems and qualitative reasoning community.

2 Introduction

Within the research community of qualitative reasoning and model-based systems, there have always been biomedical researchers who explored the ideas developed within the community in a real-world biomedical setting. This is not surprising, as understanding the mechanisms behind life is an endeavour that can only be achieved through the development of models. A considerable portion of this knowledge is essentially qualitative in nature. Similar remarks can be made with regard to clinical medicine, where diseases can only be understood and, therefore, be treated causatively, in terms of disease mechanisms as models. An overview of the processes involving biomedical models and their use within the biomedical fields is given in Figure 1.

However, the usefulness of other AI technology has also been explored in biomedicine, driven by the wish to make scientific progress. The purpose of this document is to report on synergies between model-based and qualitative approaches in artificial intelligence (AI) and other methods in AI.

There are many AI methods which have been studied within the biomedical field with regard to their applicability, such as rule-based systems with empirical associations, case-based reasoning, neural networks and so on. However, as these methods are completely different from model-based and qualitative reasoning methods, it can be doubted whether there has ever been much synergy. In this document, we therefore describe methods which are related to methods underlying model-based systems and qualitative reasoning; as a consequence, synergies can be more easily recognised.

This document is organised as follows. In Section 2 we review common ways in which causal knowledge is exploited in model-based systems using qualitative methods such as Abductive reasoning, and then describe similarities with the diagnostic use of Bayesian networks. Section 3 discusses the synergies between qualitative reasoning systems and fuzzy systems, and the role fuzzy systems can play in learning system dynamics. Inductive logic programming and qualitative reasoning are covered in Section 4. Qualitative methodologies can assist in the development of mathematical models; some work in this area is discussed in Section 5. Finally, in Section 6 the report is rounded off with some conclusions regarding expected future developments. This report complements the results

discussed in the Biomedical Road Map and the Biomedical Domain Status Document, also produced by the Biomedical Task Group.

2.1 Editorial History

Early versions of this document have been produced during a meeting of the MONET 2 Biomedical Task Group on 6th and 7th May, 2004 in Pavia, Italy, and subsequently discussed and edited at a meeting of the task group on 13th and 14th September, 2004 in Aberdeen, Scotland. The document was completed following discussions during the MONET 2 General Road Map Meeting on Capri, Italy, from 19th to 22nd October, 2004. The editorial work has been carried out by George Coghill, Liliana Ironi, Peter Lucas and René Quiniou, with final input from the Biomedical Task Group.

3 Causal Reasoning and their Quantitative Analogues

In this section we explore the relationship between qualitative approaches to model-based reasoning and quantitative approaches based on probability theory, where in particular we will look at Bayesian networks, as these allow for model-based approaches to system development exploiting causal knowledge. As Bayesian networks are often given a causal interpretation, they offer a good intermediary between qualitative reasoning methods, which are the focus of MONET2, and quantitative methods. In fact, the qualitative probabilistic network (QPN) formalism is a fully qualitative representation and reasoning method, that supports qualitative reasoning in terms of probability theory [48]. The formalism has been successfully extended during the last 8 years [41, 42].

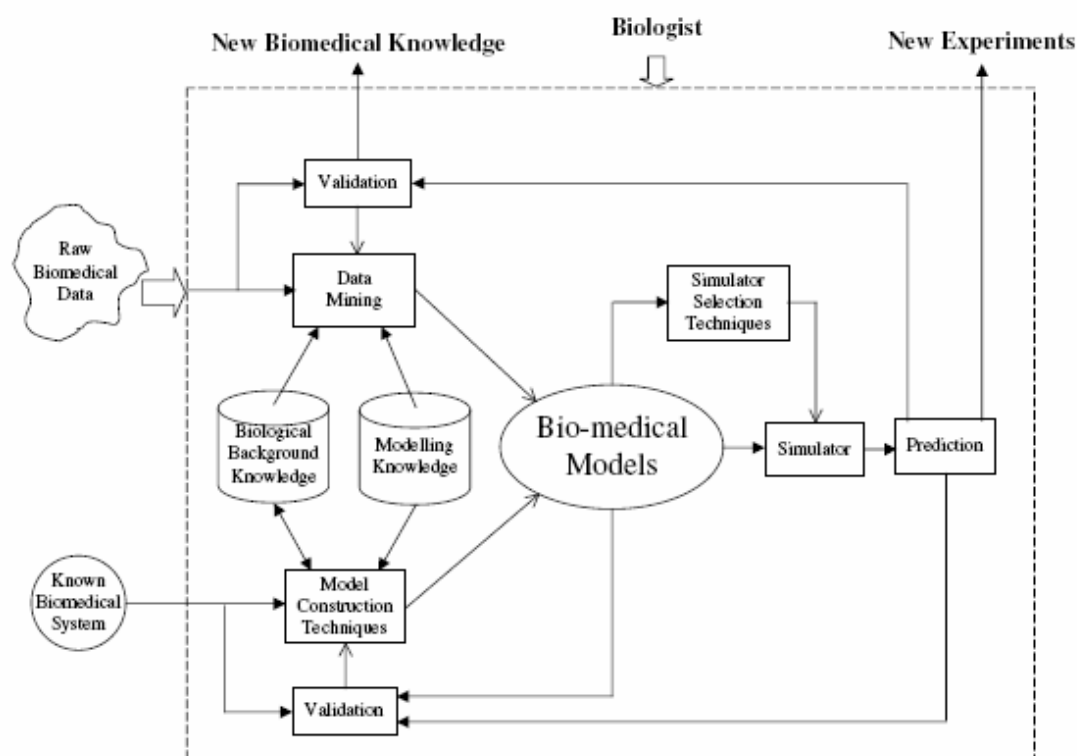


Figure 1: Overview of issues involved in the deployment of biomedical models

First, the ideas underlying causal reasoning are briefly reviewed. Subsequently, the use of causal knowledge in diagnostic reasoning and the place of abductive reasoning in this is

described. Finally, we discuss Bayesian networks and the relationship between diagnosis using qualitative approaches and diagnosis using Bayesian networks.

3.1 Causal knowledge in biomedicine

A causal model describes qualitatively a system by a set of state variables and a set of causal rules stating how these variables are affected by primary causes which either represent a normal condition or specific failure conditions. Causal rules have the general form

Conditions causes Effects

where the relationship ‘causes’ is a relation that specifies how the interaction between the elements of the collection of ‘Conditions’ give rise to the collection of ‘Effects’. For instance in the cardiac domain, we could model the effects of Left Bundle Branch Block (LBBB) as putting the left ventricle in a bad condition (here, delaying its electrical activation - DLVE) which in turn causes an enlarged QRS wave (Abn QRS) to be observed on the ECG. This knowledge is encoded by the two following rules:

LBBB causes DLVE

DLVE causes Abn QRS

LBBB is a primary cause, *DLVE* is an internal (i.e. non observable) state and *Abn QRS* is an observable state.

3.2 Abductive Reasoning and Diagnosis

Biomedical knowledge that can be captured in terms of cause-effect relationships has been extensively studied in medical AI. Diagnosis consists of inferring which primary causes could have led to a given set of observations. Intuitively, diagnostic reasoning uses the causal rules in order to relate observations to primary causes through a path of internal states. Any such path provides an explanation of a subset of observations. For instance, if an abnormal QRS is observed we could hypothesize that a delayed QRS has occurred and, further, that a left bundle branch block was the primary cause. This kind of reasoning is called Abductive reasoning. It can be contrasted to deduction reasoning which, formally, is an application of the *modus ponens* inference rule:

$$\frac{A, A \rightarrow B}{B}$$

i.e. from A and $A \rightarrow B$ deduce B . Instead, abduction uses the following inference rule:

$$\frac{B, A \rightarrow B}{A}$$

i.e. from B and $A \rightarrow B$ infer A .

It is worth noting that deduction with *modus ponens* is always correct, i.e. formulae that are deduced using *modus ponens* are also a semantic consequence using model theory. On the other hand, abduction can be incorrect: if the additional rule $C \rightarrow B$ is also available C can also be inferred abductively and C could well be the only true conclusion. Note that by here ‘incorrect’ is not meant incorrect in terms of model theory, but empirically, as abduction is

incorrect from a model-theoretical viewpoint. In effect, an abductive interpretation of causal knowledge, formalised, for example, in logical formulae, is an interpretation taking into account observed findings which are explained in terms of causal knowledge. To define the correctness of the rule of abduction it is, therefore, mandatory to incorporate notions such as explanation and observation into the interpretation of the formulae.

Abduction has been used to formalise diagnostic reasoning in the following setting [16, 30, 38, 39]. A causal model is a triplet (A, O, R) , where A is a set of propositions called the *abductibles*, O is a set of propositions called the *observations* and R is a set of causal rules having the form

Causes \rightarrow *Effects*

Given a set of observations O' , abductive diagnosis consists in applying the abductive inference rule whenever it is possible until a fixpoint is reached (saturation). The logical formalization of diagnosis is therefore:

$$R \cup H \models O'$$

where again R stands for a logical representation of a causal model and $H \subseteq A$ is the hypothesis. A common constraint on H is that it is subset minimal.

Additional logical constraints are commonly added by means of meta-level conditions, e.g.

$$R \cup H \cup C \not\models \perp$$

states that diagnostic solutions H may not be inconsistent with logical constraints C given available causal knowledge R . Thus a deduction step is performed after abduction in order to deduce all the facts that should be observed when the abducted primary causes are taken as hypotheses. If it assumed that a complete set of observations is available then not observing some deduced fact indicates that the primary causes implying this fact must be ruled out.

The conclusions obtained by abduction must be considered as being assumptions consistent with some set of observations. In some applications there can be many diagnoses. So, sometimes, integrity constraints, stating for instance that two states are mutually exclusive, can be added to rule out some diagnoses. When modelling dynamic systems, another way to refine the diagnosis is to introduce temporal constraints restricting the delay between the occurrence of the cause and the occurrence of the effects [10]. A temporal causal rule has the form $A \rightarrow B[t_B^-, t_B^+] \wedge C[t_C^-, t_C^+]$ which means that once A has occurred B will occur in a delay d_B after the occurrence of A , with $t_B^- \leq d_B \leq t_B^+$, and C will occur in a delay d_C , with $t_C^- \leq d_C \leq t_C^+$. In this setting the abduction process remains logically the same but it is now associated with constraint propagation ensuring that temporal aspects are satisfied. These temporal constraints are modelled after the physical properties of the system under study.

Generally, observations come from sensors placed on the system to be diagnosed. Since causal models are essentially qualitative, raw data coming from sensors has to be abstracted before being used in the causal model. This may involve aggregating measures from different sensors or integrating measures during some periods of time. As non linear relationships may be concerned, tools such as artificial neural networks are particularly qualified for this task. This is what was achieved in Calicot [9] where the ECG is first abstracted as a series of timestamped QRS events which are further qualified as normal or

abnormal. Signal processing algorithms are used to detect the occurrences these particular waves on the ECG, then an ANN determines their shapes. The ANN was trained on a set QRS waves extracted from an ECG database. A complete set of examples could also have been generated from simulations of a cardiac model such as Carmen [23, 24]. Carmen is a semi-qualitative model which models the physical and electrical properties of cardiac cells as cellular automata.

In addition to logic, many other representation formats have been used to represent causal knowledge for the purpose of diagnosis, such as set theory [37, 30, 31]. In terms of set theory, a diagnostic problem can be formalised as follows:

$$e(H) \supseteq O'$$

where again O stands for a set of findings that have been observed, H for a potential diagnostic explanation in terms of the function e , which associates observable effects to sets of disorders. Often, some additional constraints are formulated with respect to H , such as subset minimality.

3.3 Bayesian networks for model-based diagnosis

Although Bayesian networks can be looked upon as representing unique factorisations of joint probability distributions taking into account statistical (in)dependence constraints, they are often given a causal reading. This guides and facilitates the development of Bayesian networks models as in particular biomedical knowledge, which normally includes at least some knowledge concerning (patho)physiology, which can quite easily be captured in terms of cause-effect relationships. It is interesting to note that Bayesian networks can also be given a completely qualitative interpretation. The underlying theory will be briefly reviewed below. As a consequence of the causal reading that can be given to Bayesian networks, there is also a natural relationship between using Bayesian networks for diagnostic purposes and diagnosis using abductive reasoning. The latter is called *abductive diagnosis* in the following.

Diagnosis using explicit models of structure and functions, also called first-principle models, has been investigated extensively in the field of model-based reasoning using qualitative methods, but only on a very limited scale in the context of Bayesian networks [33]. Much of what follows will therefore focus on the relationship between abductive diagnosis and establishing a diagnosis using Bayesian networks.

3.3.1 Qualitative probabilistic networks

Qualitative probabilistic networks, or QPNs for short, are qualitative abstractions of Bayesian networks, bearing a strong resemblance to their quantitative counterparts [48]. A qualitative probabilistic network equally comprises a graphical representation of the interdependences between statistical variables, once again taking the form of an acyclic digraph. Instead of conditional probabilities, however, a qualitative probabilistic network associates signs with its digraph. These signs serve to capture the probabilistic influences and synergies between variables.

A *qualitative probabilistic influence* between two variables expresses how the values of one variable influence the probabilities of the values of the other variable. For example, a *positive qualitative influence* of a variable A on its effect B , denoted $S_+(A,B)$, expresses that observing the value \top for A makes the value \top for B more likely, regardless of any other direct influences on B , that is,

$$\Pr(b \mid a, x) \geq \Pr(b \mid \bar{a}, x) \quad (1)$$

for any combination of values x for the set $\pi(B) \setminus \{A\}$ of causes of B other than A . A *negative qualitative influence*, denoted $S^-(A, B)$, and a *zero qualitative influence*, denoted $S^0(A, B)$, are defined analogously, replacing \geq in the above formula by \leq and $=$, respectively. If the influence of A on B is non-monotonic, that is, the sign of the influence depends upon the values of other causes of B , or unknown, we say that the influence is *ambiguous*, denoted $S^? (A, B)$. With each arc in a qualitative network's digraph an influence is associated.

The set of influences of a qualitative probabilistic network exhibits various convenient properties [48]. The property of symmetry guarantees that, if the network includes the qualitative influence $S^+(A, B)$, then it also includes $S^+(B, A)$. The property of *transitivity* asserts that the qualitative influences along a path between two variables, specifying at most one incoming arc for each variable, combine into a single compound influence between these variables with the \otimes -operator from Table 1. The property of *composition* further asserts that multiple qualitative influences between two variables along parallel paths combine into a compound influence between these variables with the \otimes -operator. In addition to influences,

Table 1: The operators for combining signs.

| \otimes | + | - | 0 | ? | \oplus | + | - | 0 | ? |
|-----------|---|---|---|---|----------|---|---|---|---|
| + | + | - | 0 | ? | + | + | ? | + | ? |
| - | - | + | 0 | ? | - | ? | - | - | ? |
| 0 | 0 | 0 | 0 | 0 | 0 | + | - | 0 | ? |
| ? | ? | ? | 0 | ? | ? | ? | ? | ? | ? |

a qualitative probabilistic network includes *synergies* modelling interactions between influences. An *additive synergy* between three variables expresses how the values of two variables jointly influence the probabilities of the values of the third variable. For example, a *positive additive synergy* of the variables A and B on their common effect C , denoted $Y^+(\{A, B\}, C)$, expresses that the joint influence of A and B on C is greater than the sum of their separate influences, regardless of any other influences on C , that is,

$$\Pr(c \mid a, b, x) + \Pr(c \mid \bar{a}, \bar{b}, x) \geq \Pr(c \mid a, \bar{b}, x) + \Pr(c \mid \bar{a}, b, x) \quad (2)$$

for any combination of values x for the set of causes of C other than A and B . *Negative*, *zero*, and *ambiguous additive synergy* are defined analogously. A qualitative network specifies an additive synergy for each pair of causes and their common effect in its digraph.

A *product synergy* between three variables expresses how the value of one variable influences the probabilities of the values of another variable in view of an observed value for the third variable. For example, a *negative product synergy* of a variable A on a variable B given the value \top for their common effect C , denoted $X^-(\{A, B\}, c)$, expresses that, given c , the value \top for A renders the value \top for B less likely, that is,

$$\Pr(c \mid a, b, x) \cdot \Pr(c \mid \bar{a}, \bar{b}, x) \leq \Pr(c \mid a, \bar{b}, x) \cdot \Pr(c \mid \bar{a}, b, x) \quad (3)$$

for any combination of values x for the set of causes of C other than A and B . *Positive*, *zero*, and *ambiguous product synergy* again are defined analogously. For each pair of causes and their common effect, a qualitative probabilistic network specifies two product synergies, one for each value of the effect. Upon observation of a specific value for a common effect of two causes, the associated product synergy induces an influence between the two causes; the sign of this influence equals the sign of the synergy. A qualitative influence that is thus induced by a product synergy is termed an *intercausal influence*.

As discussed above, qualitative probabilistic networks can be used for uncertainty reasoning in a completely qualitative fashion. However, they have also been used to provide a semantic underpinning of standard probability distributions constituting Bayesian networks [34]. It is thought that such qualitative description could be useful in analysing sparse biomedical data, in particular concerning few samples, as in bioinformatics and clinical medicine.

3.3.2 Bayesian networks and maximum a priori assignment diagnosis

The connection between multiple-disorder diagnosis in approaches to qualitative model-based diagnosis as described above and probability theory is explored in this section.

Establishing a diagnosis for an individual patient in essence amounts to constructing a hypothesis about the disease the patient is suffering from, based upon a set of indirect observations from diagnostic tests. Diagnostic tests, however, generally do not serve to unambiguously reveal the condition of a patient: the tests typically have true-positive rates and true-negative rates unequal to 100%. To avoid misdiagnosis, the uncertainty in the test results obtained for a patient should be taken into consideration upon constructing a diagnostic hypothesis. Bayesian networks offer a natural basis for this type of reasoning with uncertainty. A significant number of network-based systems for medical diagnosis have in fact been developed in the past and are currently being developed. Well-known early examples are the Pathfinder [21, 22] and Munin [2] systems; a more modern system, which also includes connections to decision theory is the PTA (Pneumonia Therapy Advisor), which includes a diagnostic component [32].

Formally, a diagnosis may be defined as a value assignment \mathcal{D}^* to a subset of the random variables concerned, such that

$$\mathcal{D}^* = \arg \max_{\mathcal{D}} \Pr(\mathcal{D} \mid \mathcal{E})$$

where \mathcal{E} is the observed evidence, composed of symptoms, signs and test results. A diagnosis thus is a maximum a posteriori assignment (MPA), to a given subset of variables. Establishing a MPA from a Bayesian network, however, is extremely hard from a computational point of view. Since in addition combinations of disease do not occur very often, diagnostic reasoning is generally focused on single diseases. One approach is to assume that all diseases are mutually exclusive. The different possible diseases then are taken as the values of a single disease variable. This is the assumption underlying Pathfinder [21, 22] and many other diagnostic systems based upon probability theory. Another approach is to capture each possible disease by a separate variable. Reasoning then amounts to computing the probability distribution for each such variable separately, which is the marginal probability distribution for each of the variables V given the evidence \mathcal{E} , i.e. $\Pr(V \mid \mathcal{E})$. The combination of the most likely values for these separate disease variables, however, need not be a MPA to these variables. Nevertheless, this approach is the standard one on Bayesian networks at the moment. As an example, consider the simple Bayesian network shown in Figure 2. The structure of the graph of the Bayesian networks

indicates that the variables flu and pneumonia are (statistically) independent, as there is no link between flu and pneumonia. This implies that it holds that

$$\Pr(\text{FLU}, \text{PNEUMONIA}) = \Pr(\text{FLU}) \Pr(\text{PNEUMONIA})$$

However, after entering evidence concerning body temperature the variables FLU and PNEUMONIA have become dependent, and so it does *not* hold that

$$\Pr(\text{FLU}, \text{PNEUMONIA} \mid \text{TEMP} > 37.5) = \Pr(\text{FLU} \mid \text{TEMP} > 37.5) \Pr(\text{PNEUMONIA} \mid \text{TEMP} > 37.5)$$

Yet, only the probability distributions

$$\Pr(\text{FLU} \mid \text{TEMP} > 37.5)$$

And

$$\Pr(\text{PNEUMONIA} \mid \text{TEMP} > 37.5)$$

are shown in the diagram in Figure 2, as is common in the field of Bayesian networks; the probability distribution

$$\Pr(\text{FLU}, \text{PNEUMONIA} \mid \text{TEMP} > 37.5)$$

has not been computed, even though it could have been computed in principle (ignoring computational difficulties, which, in fact, do not hold for this case given the singly-connected structure of the graph). This illustrates the fact that whereas the model-based systems and qualitative reasoning community has put much emphasis on the issue of multiple-disorder diagnosis and did so almost from the beginning of the field, up until now Bayesian network researchers have taken this issue less seriously, even though this may give rise to strange interpretations of the results obtained by Bayesian networks.

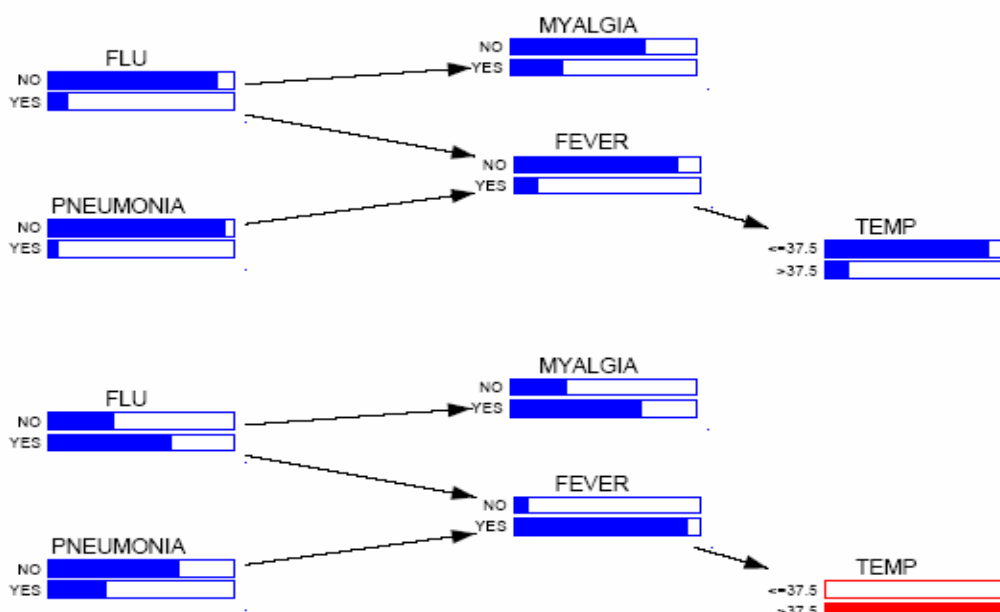


Figure 2: Marginal probability distributions of a simple Bayesian network before and after evidence concerning body temperature has been entered. Note the changes in the probability distribution of flu and pneumonia, and also note that the probability of the presence of myalgia (muscle pain) has increased. This is an example of intercausal reasoning.

To assist physicians in the complex task of diagnostic reasoning, a Bayesian network is often equipped with a test-selection method that serves to indicate which tests had best be ordered to decrease the uncertainty about the disease present in a specific patient [1]. A testselection method typically employs an information-theoretic measure for assessing diagnostic uncertainty. Such a measure is defined on a probability distribution over a disease variable and expresses the expected amount of information required to establish the value of this variable with certainty. An example measure often used for this purpose is the Shannon entropy. The measure can be extended to include information about the costs involved in performing a specific test and about the side effects it can have. Since it is computationally hard to look beyond the immediate next diagnostic test, test selection is generally carried out non-myopically, that is, in a sequential manner. The method then suggests a test to be performed and awaits the user's input; after taking the test's result into account, the method suggests a subsequent test, and so on.

3.4 The relationship between qualitative and Bayesian approaches

Above we have already touched upon the synergies between model-based systems and qualitative reasoning, on the one hand, and Bayesian networks, on the other hand. For example, we have seen that there are purely qualitative ways to reason with uncertainty, and that a qualitative notion of causality is a natural way to interpret Bayesian networks. In the context of model-based diagnosis we have seen that notions such as multiple-disorder diagnosis as proposed by the model-based systems and qualitative reasoning community has also been adopted by the Bayesian network community, although the notion is less dominant in the latter field. In this section, we discuss some similarities between model-based diagnosis using abductive reasoning as proposed by the model-based systems community and diagnostic reasoning using Bayesian networks.

A minor extension to the standard way in which model-based diagnosis is formalised within the model-based systems community is obtained by attaching a number to a diagnosis D , indicating some sort of likelihood. This is something that was already proposed by the model-based systems community many years ago. For example, in [11], [44] and [45] the standard notion of diagnosis is extended to the notion of minimal-cost diagnosis. A diagnosis D of a set of observed findings O is called a *minimal-cost explanation* of D if

$$\sum_{d \in D} \text{cost}(d) \leq \sum_{d \in D'} \text{cost}(d)$$

for each diagnosis D' of O , where cost is a function associating real values with disorders (defects) $d \in D$. The cost of a diagnosis may be anything, varying from financial costs to some subjective feeling of importance expressed by numbers. However, Charniak and Shimony (cf. [11]) choose as a semantics of cost function information for the negative logarithm of probabilities, i.e.

$$\text{cost}(D) = -\log \Pr(D)$$

Clearly, under this interpretation, a minimal-cost diagnosis is identical to a most probable diagnosis (MPA diagnosis as discussed above), as is proved by Charniak and Shimony in [11].

4 Qualitative Reasoning and Fuzzy Logic

In this section we will describe the synergies achieved between constraint based qualitative reasoning and fuzzy logic on the one hand to create a fuzzy qualitative reasoning system and on the other hand, to utilise *a priori* structural knowledge gleaned from a qualitative

model to improve the robustness and interpretability of non-linear dynamic fuzzy systems learned from input/output data.

4.1 Background

Qualitative simulation consists of a set of eclectic, theoretical, techniques designed to generate an object level qualitative description of the behaviour of physical systems from an explicit description of its **structure**, a set of initial conditions and some initial ‘disturbance’. The term ‘qualitative’ has been used in many ways, and generally to mean ‘non-numerical’ models. However, the whole field of Artificial Intelligence utilises non-numeric symbolic models. We will, therefore, use the term ‘qualitative modelling’ to mean reasoning about systems characterized by continuously changing variables (of time) by discrete abstractions of the value of the associated variables. Most of the work on qualitative modelling concerns identifying appropriate ‘abstractions’ that allow the important distinctions or landmarks, in the behaviour to be computed. This requires quantisation of the real-number line into a finite set of distinctions, known as the quantity space; one of the main goals of qualitative modelling research is to identify useful discretisations, or precisions, for particular generic purposes.

Abstract descriptions of state make it possible to have more concise representations of behaviour. However, the generation of the behaviour from such descriptions tends not to produce a unique solution. This, of course, is to be expected, as the information required to produce a unique description has been eliminated in the intentional abstraction. Therefore, with respect to the real-valued description, qualitative models may produce ambiguous descriptions of the behaviour. However, such ambiguous behaviour contains useful information. For example, if it is required to predict whether the current state can lead to a critical or dangerous condition, it may be sufficient to show that none of the possible behaviours leads to a critical situation. It is important to show, therefore, that the set of possible behaviours includes the ‘actual’ behaviour of the system i.e. that the inferences made are sound. In this way, a purpose can be satisfied even with incomplete descriptions. Whereas, in traditional methods, all of the information needs to be available, and it needs to be precisely and uniquely characterised, before any inference can be made.

Dynamic system simulation is an increasingly important form of qualitative reasoning. The pioneer of this approach has been the QSIM algorithm of Ben Kuipers [29]. This approach epitomises the constraint ontology, and utilises sets of qualitative equations with a landmark quantity space. One of the attractive features of QSIM is that it is designed to handle incompleteness in the knowledge of the model. The incompleteness here takes the form of a lack of knowledge concerning functional relations in the system. This situation is captured by the monotonic function constraints M^+ and M^- between two variables, which declares that one variable monotonically increases (+) or decreases (-) with respect to another variable, covering families of relations.

The simulation algorithm is a particular form of the above generate-and-test procedure from the current state assuming only that the variables are continuous. These successor states are then checked for consistency with the modelling constraints and conflicting states are eliminated. Thereby directionality of the constraints is not required and hence no causal interpretation of the algebraic constraints of the model can be derived. The algorithm can, however, be shown to predict the ‘correct’ behaviour, interpreted as the qualitative abstraction of the real-valued solution of the associated differential equation, thereby associating the causality with the integration part of the process as captured in the transition rules of the simulation engine.

On the other hand, one of the well recognised aspects of human reasoning is that it is inherently vague. It appears that the use of such terms as “quite high” and “usually” are

essential to coping with the uncertainties and complexities found in real-world situations. Any formalism purporting to capture such knowledge must, therefore, support the representation of and reasoning with such vague and imprecise knowledge. One way to represent vagueness is by the use of Fuzzy Sets [50]. These were proposed in the middle of the 1960's by Lotfi Zadeh, and after an early flourish and a quiet 1980's have become very popular in the 1990's, due to the adoption of the techniques by the Japanese within domestic products.

A complete theory of fuzzy logic has been developed. This allows the negation, conjunction and disjunction of fuzzy sets through the use of max and min operators that can be seen as a generalisation of classical crisp logical connectives. Fuzzy reasoning can also be supported by the compositional rule of inference that allows sets in different universes of discourse to be related through a set of fuzzy relations. This forms the basis of a mechanism for fuzzy rule-based systems where antecedents are represented by discrete fuzzy sets and the strengths of relationships, determined from fuzzy rules, are used to produce fuzzy consequents. This mechanism is, in fact, very easy to implement, using a combination of maximum and minimum of vectors of real-numbers. Indeed VLSI devices are now available that are designed specifically for fuzzy reasoning and many software systems are now commercially available that support the development of fuzzy rule-based systems. Many biomedical applications of fuzzy logic feed-back control systems are known, in fact, this was the first real application of fuzzy sets [50]. Now they are used for pattern matching, interface design and for the simulation of systems.




Figure 3: A single compartment system

The distinct advantage of fuzzy rule-based systems is that they are an explicit attempt to represent human expertise and are easily implemented. They may be considered as black box models; that is, once the fuzzy values have been adequately tuned the input / output relations can provide a good description of the behaviour of the represented system. The disadvantage is that they cannot provide a justification for their results because they lack structure; if they work they appear to work well, but it is difficult to establish the limits of their operability. The properties of fuzzy rule-based systems differ from those of the rule-based systems only in the representation of precision and uncertainty.

4.2 Synergies between Qualitative Models and Fuzzy Systems

4.2.1 Fuzzy Qualitative Simulation

In order to achieve the explanatory power of structured qualitative models combined with the advantage of increased precision in the constitutive relations of the qualitative model, Fuzzy Qualitative Systems were developed [46, 12, 13]. Fuzzy qualitative models are structurally the same as qualitative models. However, they differ in that the variables of the model take their values from a fuzzy quantity space and the M^+ and M^- relations are replaced by fuzzy rule bases. That is: a fuzzy rule based system forms a part of the structure of a qualitative model.

One example of a fuzzy qualitative simulator is the *Morven* toolset. This is a constraint-based fuzzy qualitative reasoning system which built on the seminal work of Shen & Leitch; it contains a number of simulation and envisionment algorithms. The development of this toolset has permitted the suitability of different techniques to be examined in a number of contexts; and the comparison of different approaches to constraint-based fuzzy qualitative simulation to be made [12].

As an example consider the case of Fuzzy Vector Envisionment. This may be illustrated by means of the simple example of a single compartment system as shown in figure 3.

This system is first order and can be represented by a first order differential equation of the form:

$$x'(t) = u(t) - k_{10} \cdot x(t)$$

from which a qualitative differential equation may be abstracted. Here $x(t)$ is the concentration of substance in the compartment, $u(t)$ is the infusion into the compartment and k_{10} the rate of elimination from the compartment

If one were to model this system purely qualitatively, the situation where the rate of infusion were constant could only be given by a value of 'positive' or '+' for $u(t)$. This situation yields very useful results in the form of a global picture of the possible behaviours of the system at the lowest precision. The integration of fuzzy sets with qualitative simulation enables the generation of envisionments at higher levels of precision. The results of one such are shown in figure 4 where the complete envisionment for a steady infusion to the compartment is the union of the envisionment graphs for each possible value of the magnitude of $u(t)$: p - small, p - medium, p - large or p - max in this case.




Figure 4: A fuzzy envisionment of a single compartment system

This approach to the merging of fuzzy and qualitative techniques captures strengths from both, as shown above, and permits the incorporation of temporal and empirical reasoning within a qualitative simulation in an integrated manner.

3.2.2 Learning System Dynamics through Fuzzy Models

In the fuzzy modelling approach, the identification of the nonlinear system dynamics deals with the reconstruction of functional relationships $f: X \rightarrow Y^1$ between the input-output variables from the available data samples only. The function $f(\cdot)$ is unknown, and the modelling problem is tackled by looking for a continuous function approximator \tilde{f} of f . In this modelling framework, the functional form of \tilde{f} results from the mathematical interpretation of the fuzzy operators of M fuzzy rules: such rules are expressed through if-then rules, and describe the knowledge on the relations between the input-output variables. The components of model parameter vector θ , which are tuned on the experimental data through optimization procedures, are associated with the membership functions of input-output variables or, in other words, with the locations of their fuzzy partition. Therefore, the modelling problem deals with the determination of the fuzzy partition of the input-output variables, and the optimal number of rules that must be used to generate the fuzzy model. In theory, both partitions and inference rules can be derived by the expert knowledge, but such information may be very poor, irregular and unstructured, and then, in practice, prevents one from defining the optimal form of $\tilde{f}(\cdot)$, where by optimal we mean that $\tilde{f}(\cdot)$ is of minimal complexity, but able to capture all of the significant features of the system dynamics. For these reasons, the research efforts turned to the definition of learning methods that automatically generate the fuzzy systems from the data samples only [47]. Although these methods have been successfully applied to a variety of domains, they are affected by two serious drawbacks. Let us observe that the same problems may occur even when the resulting model is abstracted from the expert knowledge, since an empirical rather than structural kind of knowledge is mostly given.

A part serious inefficiency problems due to the possible combinatorial explosion of the number of rules and parameters when the number of inputs increases, these methods suffer from lack of robustness and interpretability. As for robustness, the following problems may occur:

1. generalization capability of the built model may be lost as overfitting phenomena may occur, especially when the data set is not large enough as it may often occur in the biomedical domain;
2. numerical instability phenomena during parameter estimation, e.g. solutions do not necessarily depend on the data in a continuous way, may occur.

In addition, fuzzy models generated from data only do not guarantee to gain insight into the system: the resulting model structure is often not transparent, and, after their optimization, the model parameters, that are associated with the locations of the fuzzy sets rather than with the biomedical reality, may lead to an incomplete, inconsistent and even indistinguishable fuzzy partition.

¹ $X \subseteq \mathbb{R}^n$ input space, and $Y \subseteq \mathbb{R}$ output space

Integration of Qualitative Models and Fuzzy Systems: a hybrid method

Let us observe that for a great deal of pathophysiological systems, the available structural knowledge is insufficient for the formulation of a quantitative differential model, but does not prevent the formulation of a qualitative one. This consideration motivated the definition of a new hybrid approach, called FS-QM, to the fuzzy identification of dynamical systems [5,6]. Its novelty consists in the way the fuzzy model is built: both the fuzzy partition and rule base are defined upon the available structural knowledge. FS-QM is applicable whenever the incompleteness of a priori knowledge is such that it allows us (i) to write a QSIM model [29], and (ii) to bound the uncertainty on landmark values to a range of numerical values. In outline, the whole range of possible system dynamics, represented and simulated within the qsim modelling framework, is automatically translated into the fuzzy formalism. The domain of each input/output variable is automatically partitioned into fuzzy sets in accordance with its associated quantity space, and with the prior information on landmark numerical bounds. In other words, the cardinality of the fuzzy partition of a variable and the membership function locations are defined by the cardinality of the set of qualitative values the variable may assume, and by the interval bounds of its landmarks, respectively. Given a landmark-based fuzzy partition and the simulated behaviour tree, the generation of the fuzzy rule base is straightforwardly derived by mapping each behaviour of the input/output variables into a set of rules, where each rule describes a transition from a qualitative state to the next one.

The mathematical interpretation of the rule base explicitly initializes $\tilde{f}(\cdot)$ which is then refined through parameter identification from the data samples. Fig. 5 sketches the main steps of FS-QM.

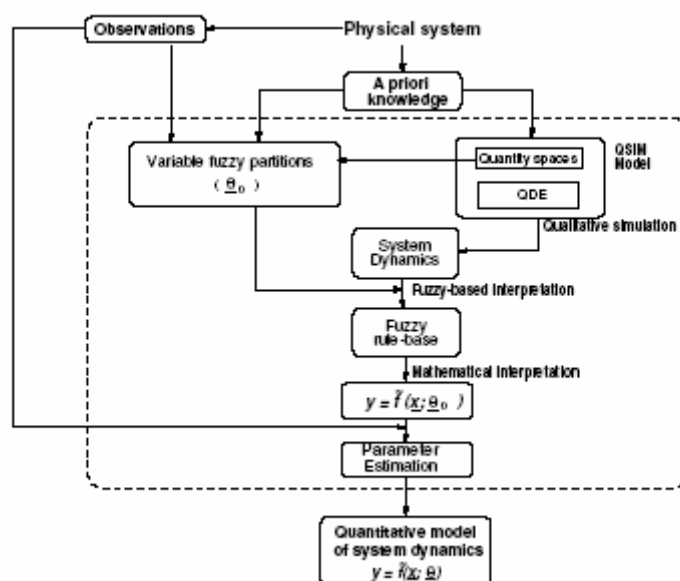


Figure 5: Main steps in FS-QM

Efficiency, robustness, and interpretability of FS-QM models

The hybrid method outlined above, where the structure identification phase is clearly separated from the parameter optimization one, allows us to get a nonlinear functional relationship between input and output variables that is both robust and interpretable [20]. It allows us to get (i) good generalization properties as the fuzzy system does represent the system dynamics, (ii) numerical stability as, due to the a priori knowledge and the particular way we perform variable partitions, regularization techniques can be applied, and,

consequently, the search of optimal parameters is restricted within a "trust" region. The regularized problem also allows us to get a fuzzy partition which is complete, consistent, and distinguishable after the parameters have been tuned on the data. This together with the fact that the parameters do refer to the reality makes the identified model completely interpretable. Moreover, the method is efficient as the number of both rules and parameters is linear with the cardinality of the fuzzy partition or, in other words, by the cardinality of the quantity space.

Application perspectives in the biomedical domain

To achieve a robust and interpretable fuzzy model, FS-QM effectively employs all the available structural prior knowledge, represented in QSIM, and empirical data. The embedding of deep prior knowledge into the function $\tilde{f}(\cdot)$ makes the identification problem better posed, since it properly delimits the model search space. In addition, the prior knowledge allows us to define a good initial estimate of the parameter vector $\underline{\theta}_0$, and then, to define a trust region where $\underline{\theta}$ is supposed to belong. If the prior knowledge is correct, this will lead to a model that has good generalization and interpretability properties even in data-poor contexts. Although FS-QM rests on correct prior knowledge expressed as a QSIM model, and in particular on physically significant predefined landmarks, its range of applicability in the domains of Life Sciences in which we are interested is quite wide, as such knowledge is available for a great deal of dynamical systems. In particular, it has already been successfully applied to model:

1. a predictor of blood glucose level in diabetic patients [4]
2. a simulator of the intracellular thiamine (vitamin B₁) kinetics in the intestine tissue [7].

Let us observe that as for the metabolic system (2), FS-QM has allowed us to solve serious identification problems: (i) the classical differential approach turned out to be inapplicable for the incompleteness of the available knowledge and for the difficulty of gathering an adequate number of experimental data;(ii) the conventional fuzzy approaches failed for the too small number of available input-output samples [7].

In FS-QM, the gained parameter interpretability from the physical point of view represents an added value that makes it possible for fuzzy models to be used to perform a larger spectrum of tasks than the usual one, namely system control. FS-QM models might be conveniently applied, for example, in a diagnostic context. On the one hand, diagnostic hypotheses that explain the observed behaviours could be tested by introducing structural variations into the underlying qualitative model, and by validating the fuzzy model built on the basis of the newly generated rule base. On the other hand, diagnostic hypotheses could be drawn from the analysis of the deviations of the estimated values of parameters from the nominal ones.

5 Qualitative Reasoning and Inductive Logic Programming

The development of intelligent tools to aid in the process of Scientific Discovery, particularly in the construction of explanatory models, is an important goal of AI; and qualitative modelling provides an ideal representation. This is the ultimate in adaption, and a hybrid system merging Inductive Logic Programming and Qualitative Simulation is a suitable tool for achieving it. Bioinformatics is an ideal domain for applying this technology: the data are sparse (making it unsuitable for numerical techniques), they are noisy and they require the construction of models which will inevitably include unobserved variables. Work on constructing models of systems in molecular biology is in the early stages of development

and so, given the above stated challenges, any useful results emerging will be of tremendous practical value.

The ultimate goal in this scientific quest is the production of quantitative models; however, the discovery of suitable structural models (qualitative differential equations) can be the means of directing the scientist as to which experiments to carry out next in the path towards this goal. In this section we present a learning system which combines Inductive Logic Programming (ILP) with QSIM in order to construct qualitative models of physical and biological systems containing unmeasured variables [14].

The general model learning problem can be represented deductively as follows: if we term the observations (evidence) E , the background knowledge B , and the hypothesis to be learnt H , then given that:

$$B \not\models E \quad (4)$$

find a hypothesis H such that

$$B \wedge H \models E \quad (5)$$

Many solutions to this problem are possible, e.g. the trivial solutions of E , or $B \rightarrow E$. The problem is therefore how to restrict solutions to suitable ones. In abduction [18] solutions are restricted to ground facts; in ILP more general solutions are allowed [35], although there are still typically syntactic restrictions on what form solutions can take. For most scientific discovery problems it is clear that ILP is advantageous, as we wish to learn general theories; and for similar reasons ILP is a sensible choice for learning QSIM models.

ILP is distinguished from other machine learning techniques by using first-order predicate logic (specifically logic programs) to represent background knowledge, observations, and hypotheses [36]; and we have previously applied machine learning and ILP to many scientific problems with success (e.g. [28]).

As discussed in the previous section QSIM [29] is a constraint based qualitative simulation engine and utilises an equational representation which is an abstraction of *ordinary differential equations*. It is the most highly developed constraint based Qualitative Reasoning (QR) system available.

There are several kinds of constraint which can appear in a QSIM model. There are predicates, implemented as relations, representing the usual algebraic operations of addition, multiplication, and sign inversion; plus a derivative predicate stating that one variable is the derivative of another. The conjunction of qualitative relations models the relationships between a set of measured variables, plus a number of unmeasured variables. There may be zero, one or more unmeasured variables, which we term the model's *hidden variables*. Where there are sufficient hidden variables, the method described here can discover *hidden relations* that relate only hidden variables; this is a novel feature of the learning system presented here. The learning of qualitative models from examples is a great challenge for current machine learning methods since the search space is very large. The problem is also interesting because the data are *positive only*, i.e. when identifying a system, nature only provides positive examples of states of the system, not examples the system can *not* be in. This hinders machine learning as there are no negative examples to restrict over-generalisation.

It is easiest to explain the operation of this system by means of an illustration. Consider a single compartment system with a single input infusion (constant). Fig. 6 shows the physical system and the corresponding qualitative representation.

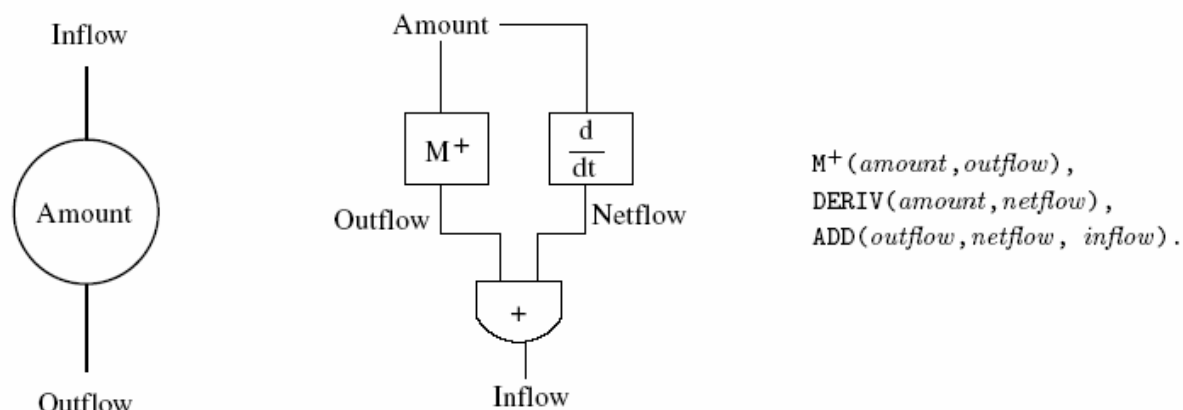


Figure 6: The single Compartment system (left) physical; (middle) diagrammatic; and (right) QSIM relations. Here, inflow, outflow and amount have the obvious meanings and are the measured variables. “Netflow”, the net difference between the inflow and outflow, is an intermediate (unmeasured) variable.

The compartment is initially empty. The infusion is turned on to give a steady in- flow to the compartment, and measurements of the inflow, outflow and the amount of substance in the compartment (actually concentration) are obtained. These are converted (by a quantitative-to-qualitative converter) into a sequence of qualitative observations (or *states*). Each observation is a tuple stating the magnitude and derivative direction for the qualitative variable:

| State | Inflow | Outflow | Amount |
|-------|-------------------------------|-------------------------------|-------------------------------|
| 0 | $\langle(0, inf), std\rangle$ | $\langle 0, inc\rangle$ | $\langle 0, inc\rangle$ |
| 1 | $\langle(0, inf), std\rangle$ | $\langle(0, inf), std\rangle$ | $\langle(0, inf), std\rangle$ |
| 2 | $\langle(0, inf), std\rangle$ | $\langle(0, inf), inc\rangle$ | $\langle(0, inf), inc\rangle$ |
| 3 | $\langle(0, inf), std\rangle$ | $\langle(0, inf), dec\rangle$ | $\langle(0, inf), dec\rangle$ |

Given these qualitative observations as examples and background knowledge consisting of constraints on models (described later) and QSIM relations Qoph performs a search for acceptable models. Assuming, for simplicity, that the only QSIM relations available to Qoph are M^+ , $DERIV$ and ADD , Fig. 7 shows the space of possible models to be searched.

activity which usually demands a long and close collaboration of many experts. QR can play a variety of synergic roles within this process:

1. it may provide methods for a preliminary qualitative processing of data to select classes of mathematical models that guarantee adherence to observed features, and even drive the initialization of numerical methods (see Subsection 5.1). Consequently, it may lead to the development of automated modelling tools that can support experimental researchers in the investigation of structural mechanisms underlying physical processes.
2. it can significantly help the process of understanding mechanisms underlying complex physical phenomena by providing methods for reasoning about spatially distributed / temporal data, as they can also result from numerical simulations (see Subsection 5.2).

Besides helping medical research in the important phase of the definition of interpretative rationales through models and their simulation, QR methods can lead to the automated interpretation of numerical fields in specific medical domains, and therefore to the realization of tools that can eventually enter the clinical practice.

6.1 QR-based automated modelling: new perspectives in pharmaceutical design

To fully accomplish the task of drug design, a variety of knowledge and skills are required of researchers. A drug basically consists of an active ingredient embedded within a carrier whose chemico-mechanical properties are fundamental to meet such requirements as bio-availability or patient's compliance. Therefore, rheological properties as well as more specific chemicophysical properties of the carrier such as mucoadhesion, hydrophilia, etc. that affect bioavailability need to be thoroughly explored. The assessment of such drug vehicle properties have traditionally been carried out experimentally by pharmaceutical researchers in spite of its high costs. The availability of automated modelling tools able to exploit all the available domain specific and domain-independent knowledge can therefore significantly help improve the cost-effectiveness ratio involved in drug design. As a matter of fact, modelling requires knowledge and skills that are not familiar to these experts. By modelling the mechanical behaviour of the polymeric visco-elastic materials, commonly used as drug vehicles, it is possible to characterize them with respect to their physico-chemical properties, shedding light on the underlying mechanisms, and providing a cheaper way to assess materials and tune the drug delivery system to meet the desired pharmacological properties.

A computational framework was developed by Rossi et al. and Ironi and Tentoni [43, 25] that fulfils the automated modelling task for the wide class of visco-elastic materials. The system takes as input data collected by performing a set of standard rheological tests on a sample of material, and provides as output a quantitative ODE model, i.e. the constitutive law of the material, whose structure and numerical parameters fully characterize the material complexity (number and strength of chemical bonds, and molecular entanglement). The system integrates different kinds of knowledge representation, modelling and simulation methodologies, making the most out of the available knowledge. Pattern recognition and QR techniques play a role in abstracting the salient qualitative features from the input data that allow to draw from a pre-built library of models a class of ODE's whose simulated qualitative behaviour is consistent with the observed behaviour. Then, numerical system identification methods are exploited to provide the output quantitative model. The library of models results from the exploitation of QR and symbolic computation techniques. The system has already been applied to the assessment of mucoadhesive materials [43].

6.2 QR-based Automated Data Interpretation: New Perspectives in Electrocardiographic Imaging

To gain a deeper comprehension of the electrical phenomena in the heart, complex partial differential equation models of the bioelectric sources have been developed, that take into account the fibre architecture and conduction anisotropy [19]. By simulating the wavefront propagation in simple situations, researchers can better understand how to relate visual features to underlying electrical events, both in normal and pathological cases. But, such interpretative skills belong only to very few domain experts.

Electrocardiology is a stimulating application domain where QR methodologies can fruitfully support traditional quantitative techniques in the investigation and comprehension of complex phenomena. Nowadays in clinical practice, information about the heart's electrical activity is routinely gathered through electrograms (ECG), which record electrical potential from just nine sites on the body surface. However, thanks to the latest technological advances, body surface potential maps are becoming available, as well as epicardial maps obtained non invasively from body surface data through mathematical model-based reconstruction methods. QR research effort has so far been mainly aimed at the automated interpretation of ECG's, for which the interpretative rationale is well established [8, 49]. However, a number of electrical conduction pathologies (arrhythmias, Wolf Parkinson White syndrome, just to cite two) may be missed by ECG analysis, and could instead be captured by electrocardiographic maps.

An important role in the process of defining an interpretative rationale for electrocardiographic maps can be played by QR methodologies for spatial/temporal reasoning that could 1) support the expert electrocardiologist in identifying salient features in the maps, and 2) achieve the long term goal of automating map interpretation to be used in a clinical context.

In this perspective, a novel QR approach, namely Spatial Aggregation (SA, [3]), is being used to identify patterns and salient features in epicardial activation isochronal maps [26]. In this kind of map (Figure 8), the time at which each point starts activating, derived from the electrical data of a whole heart beat, is visualized by means of isocurves.

A lot of information about the excitation wavefront structure and propagation is summarized in a single such map, since isocurves represent subsequent snapshots of the travelling wavefront. Breakthrough locations, high and low velocity pathways, conduction block regions, for example, are salient features that characterize the heart's electrical activity: they visually correspond to specific geometric patterns to be identified in the map, such as minima location, maximum and minimum elongation directions in the isocurves shapes.




Figure 8: Activation isochronal map generated on a model ventricular surface (cylindrical projection)

Unlike traditional visualization tools, the SA approach is specifically designed for the interpretation of numeric fields that are spatially represented. Global patterns can be identified and structural information about the underlying events can be captured through a hierarchical strategy in aggregating and abstracting geometrical objects at higher and higher levels. Such abstraction processes rely on the characterization of classes of geometric objects by means of relations which involve the physical variables and are spatially propagated by contiguity. So, for example, isochrones are abstracted from the activation time field by aggregating all and none but the contiguous points that share the same excitation state. The success of SA methods strongly relies on the neighbourhood graphs that explicate contiguity among geometric objects, as well as the equivalence relations that are propagated on them. Such definitions and algorithms are typically task dependent, and need to be carefully developed. In this regard, robustness issues of SA methods with respect to the contouring task have been discussed [27].

7 Future Directions in Synergy

The synergies identified in this document reflect the current state with respect to the interests of the Task Group members; and inevitably this means that the work is disjoint. The nature of the objects of research in the biomedical fields, i.e. living matter as cells, biochemical pathways, physiological mechanisms, requires model building. This is always done, at least partially, from a qualitative perspective, where details are filled in step by step if possible. This implies that in many areas of the biomedical field there may come a stage where quantitative approaches are being explored, although this is neither necessarily always the case, nor always possible. This report makes clear that although the various topics studied are diverse, the general approach with respect to synergies is always strikingly similar.

The role of computers in biomedical sciences and health care has drastically increased during the past few years, and is expected to increase further in the coming decades. The complexity of the subject matter studied in the life sciences renders it impossible to obtain insight into the underlying mechanisms using a completely quantitative approach. Often, data are sparse, as in microarray experiments, where there are normally only few samples for analysis available, and are noisy and incomplete. In other biomedical applications, researchers are exploring temporal and spatial information, sometimes in isolation, but in other circumstances in combination. It is this complexity that motivated researchers in model-based systems and qualitative reasoning to develop their abstractions. It is the firm belief of the Task Group members that these abstractions may contribute to handling the complexity of the fields, and some examples of these have been mentioned in this report. This may give rise to the further exploitation of the synergies between research in model-based systems and qualitative reasoning and related areas of research, in particular AI research.

8 References

- [1] Andreassen S. Planning of therapy and tests in causal probabilistic networks. *Artificial Intelligence in Medicine* 1992; 4:227–241.
- [2] Andreassen S, Woldbye M, Falck B, Andersen SK. MUNIN - A causal probabilistic network for interpretation of electromyographic findings. In: McDermott J, editor. *Proceedings of the 10th International Joint Conference on Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann; 1987. p. 366–372.

- [3] Bailey-Kellogg C., Zhao F., and Yip K.(1996). Spatial aggregation: Language and applications. In *Proc. AAAI-96*, 517-522, Los Altos:Morgan Kaufmann.
- [4] Bellazzi R., Ironi L., Guglielmann R., and Stefanelli M. (1998). Qualitative models and fuzzy systems: an integrated approach for learning from data. *Artificial Intelligence in Medicine*, 14, 5-28.
- [5] Bellazzi R., Guglielmann R., and Ironi L. (1999). A qualitative-fuzzy framework for nonlinear black-box system identification. In Dean T., editor, *IJCAI 99*, 1041-1046, San Francisco:Morgan Kaufmann.
- [6] Bellazzi R., Guglielmann R., and Ironi L. (2000). How to improve fuzzy-neural system modelling by means of qualitative simulation. *IEEE Trans. on Neural Network*, 11:1, 249-253.
- [7] Bellazzi R., Guglielmann R., Ironi L., and Patrini C. (2001). A hybrid input-output approach to model metabolic systems: An application to intracellular thiamine kinetics. *Journal of Biomedical Informatics*, 34, 221-248.
- [8] Bratko I., Mozetic I., and Lavrac N. (1989). *Kardio: A Study in Deep and Qualitative Knowledge for Expert Systems*. MIT Press, Cambridge, MA.
- [9] Carrault G., Cordier M.O., Quiniou R., Wang F. (2003). Temporal abstraction and Inductive Logic Programming for arrhythmia recognition from electrocardiograms, *Artificial Intelligence in Medicine*, 28, 231-263.
- [10] Brusoni B., Console L., Terenziani P., Theseider Dupr'e D. (1998). A spectrum of definitions for temporal model-based diagnosis. *Artificial Intelligence*, 102, 39–79.
- [11] Charniak E, Shimony SE (1994). Cost-based abduction and MAP explanation. *Artificial Intelligence*, 66, 345–374.
- [12] Coghill G.M. Morven: A framework for fuzzy qualitative reasoning Ph.D. Thesis, Heriot-Watt University, September 1996.
- [13] Coghill G. M. On the development of fuzzy envisionment. *Proceedings of the third International Workshop on Hybrid Methods and Adaptive Systems (CD)* Oulu, Finland July 2003.
- [14] Coghill G. M., Garrett S. M., King R. D. Learning qualitative metabolic models *Proceedings of the 16th European Conference on Artificial Intelligence*, Pp 445 - 449, Valencia, Spain, 22–27 August 2004.
- [15] Console L., Theseider Dupr'e D., Torasso P. (1989). A theory of diagnosis for incomplete causal models. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, 1311–1317.
- [16] Console L., Torasso P. (1991). A spectrum of logical definitions of model-based diagnosis, *Computational Intelligence*, 7(3), 133–141.
- [17] Console L., Theseider Dupr'e D., Torasso P. (1991). On the relationship between abduction and deduction, *Journal of Logic and Computation*, 1(5), 661–690.
- [18] Flach P., Kakas, A (2000). *Abduction and Induction: essays on their relationship and integration*. Springer-Verlag, Berlin.

[19] Colli Franzone P., Guerri L., and Pennacchio M. (1998). Spreading of excitation in 3-D models of the anisotropic cardiac tissue. II. Effect of geometry and fiber architecture of the ventricular wall. *Mathematical Biosciences*, 147, 131-171.

[20] Guglielmann R., Ironi L. (2004). The need for qualitative reasoning in fuzzy modelling: robustness and interpretability issues. *18th International Workshop on Qualitative Reasoning*, Evanston August 2-4, 113–120.

[21] Heckerman DE, Horvitz EJ, Nathwani BN (1992). Towards normative expert systems: part I – The Pathfinder project. *Methods of Information in Medicine*, 31, 90-105.

[22] Heckerman DE, Nathwani BN (1992). Towards normative expert systems: part II - probability-based representations for efficient knowledge acquisition and inference. *Methods of Information in Medicine*, 31, 106–116.

[23] Hernández A., Carrault G., Mora F., Bardou A (2000). Overview of CARMEN: A new dynamic quantitative cardiac model for ECG monitoring and its adaptation to observed signals. *Acta Biotheoretica*, 48(3-4), 303–322.

[24] Hernandez AI, Carrault G, Mora F, Bardou A. (2002). Model-based interpretation of cardiac beats by evolutionary algorithms: signal and model interaction. *Artificial Intelligence in Medicine*, 26(3), 211-35.

[25] Ironi L., Tentoni S. (2003a). A model-based approach to the assessment of physicochemical properties of drug delivery materials. *Computers & Chemical Engineering*, 27: 6, 803-812.

[26] Ironi L., Tentoni S. (2003b). On the problem of adjacency relations in the Spatial Aggregation approach. *Seventeenth International Workshop on Qualitative Reasoning*, Brasilia 20-22 August, 111-118.

[27] Ironi L., Tentoni S. (2003c). Towards automated electrocardiac map interpretation: an intelligent contouring tool based on Spatial Aggregation. *Lecture Notes in Computer Science*, 2810, 397-408.

[28] King R. D., Srinivasan A. The discovery of indicator variables for qsar using inductive logic programming. *Journal of Computer Aided Molecular Design*, 11, 571–580, 1997.

[29] Kuipers, B.J. (1994). *Qualitative Reasoning: modelling and simulation with incomplete knowledge*. MIT Press, Cambridge MA.

[30] Lucas P.J.F. (1997). Symbolic diagnosis and its formalisation, *The Knowledge Engineering Review*, 12(2), 109-146.

[31] Lucas P.J.F (1998). Analysis of notions of diagnosis, *Artificial Intelligence*, 105(1-2), 293–341.

[32] Lucas P.J.F., De Bruijn N., Schurink K., Hoepelman I.M. (2000). A Probabilistic and decision-theoretic approach to the management of infectious disease at the ICU. *Artificial Intelligence in Medicine*, 19(3), 251–279.

[33] Lucas P.J.F. (2001). Bayesian model-based diagnosis. *International Journal of Approximate Reasoning*, 27: 99–119.

- [34] Lucas P.J.F. (2002). Bayesian network modelling by qualitative patterns, in: F. van Harmelen (ed.), *Proceedings of the 15th European Conference on Artificial Intelligence (ECAI2002)*, IOS Press, Amsterdam, pp. 690–694.
- [35] Muggleton S., Feng C. Efficient Induction of logic programs. *Proceedings of the 1st Conference on Algorithmic Learning Theory*, Tokyo, 1990.
- [36] Muggleton S. *Inductive Logic Programming*. Academic Press, London, 1992.
- [37] Peng Y., Reggia J.A. (1990). *Abductive inference models for diagnostic problem solving*. New York: Springer-Verlag.
- [38] Poole D., Goebel R. and Aleliunas R. (1987). Theorist: a logical reasoning system for defaults and diagnosis. In *The Knowledge Frontier* (N. Cercone and G. Mc Calla, eds.), 331-352. Berlin: Springer-Verlag.
- [39] Poole D. (1988). Representing knowledge for logic-based diagnosis. In *Proceedings of the International Conference on Fifth Generation Computer Systems 1988*, 1282–1290. ICOT.
- [40] Poole D. (1989). Explanation and prediction: an architecture for default and abductive reasoning. *Computational Intelligence*, 5(2), 97–110.
- [41] Renooij, S. *Qualitative Approaches to Quantifying Probabilistic Networks*. PhD Thesis, Utrecht University, 2001.
- [42] Renooij, S. and Van der Gaag, L.C (2002). From qualitative to quantitative probabilistic networks. In: A. Darwiche, N. Friedman, editors, *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers, San Francisco, California, pp. 422–429.
- [43] Rossi S., Bonferoni C., Caramella C., Ironi L., Tentoni S. (1999). Model-based interpretation of creep profiles for the assessment of polymer-mucin interaction. *Pharmaceutical Research*, 16: 9, 1456-1463.
- [44] Santos E. Jr (1994). A linear constraint satisfaction approach to cost-based abduction. *Artificial Intelligence*, 65, 1–27.
- [45] Santos E. Jr, Santos E.S. (1996). Polynomial solvability of cost-based abduction. *Artificial Intelligence*, 86, 157–170.
- [46] Shen Q, Leitch R. R. Fuzzy Qualitative Simulation. *IEEE Transactions on Systems, Man and Cybernetics*, 23, 1038 – 1061, 1993.
- [47] Wang L.X. (1994). *Adaptive Fuzzy Systems and Control: design and stability analysis*. Englewood Cliff, NJ: Prentice-Hall.
- [48] Wellman M.P. (1990). Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 44, 257–303.
- [49] Weng F., Quiniou R., Carrault G., and Cordier M.-O. (2001). Learning structural knowledge from the ECG. In *ISMDA-2001*, 2199, 288-294, Berlin: Springer.
- [50] Zadeh L. Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Transactions on Systems, Man and Cybernetics*, 3, 28 – 44, 1973.

9 Document History

| <i>Version</i> | <i>Date</i> | <i>Changes made to document</i> | <i>Changed by</i> |
|----------------|--------------------------------|---|-------------------|
| 1.0 | 7 th May 2004 | First draft prepared | TG |
| 2.0 | 5 th November 2004 | Sections added | GMC |
| 3.0 | 10 th November 2004 | Proof reading and minor amendments made | JNT / RIR |
| | | | |