monet Newsletter

Issue 3 Supported by EC Research Framework Programme 5 under the IST Programme ISSN 1464-9276

Monet Executives on top in South America

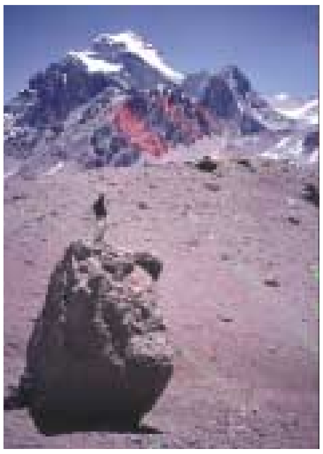

Contents

Last few Summer School places (see centre pages)

Two members of the MONET Executive committee welcomed in the New Year by climbing Aconcagua in Argentina, January 2003. Capping the Andes range at 6,869 meters (22,800 ft), Aconcagua is the highest mountain in South America and is also the highest mountain outside the Himalayas.

The ascent was by Sermatech Intelligent Applications Managing Director, Rob Milne from the UK and Louise Travé-Massuyès of LAAS-CNRS of France. The ascent of the mountain took two weeks with the round trip walk of about 115 kilometers. Base Camp was at 4,000m and two high camps were used. The highest camp was roughly the same height as the top of Kilimanjaro. Rob said, "It is like sleeping on top of Kilimanjaro (the highest mountain in Africa), and then climbing Ben Nevis (the highest mountain in the UK), but doing it five vertical miles above sea level. The biggest problem is the high winds (even worse than Scotland) and the cold temperatures. Between Base Camp and Camp 1, it was necessary to cross several fields of Penitentes. This is a rare snow pattern where the snow melts into large pinnacles. Many of them were taller than Louise." (see photo above).

Dr. Louise Travé-Massuyès is a Research Director of the *Centre National de Recherche Scientifique (CNRS),* working at *LAAS*, Toulouse, France, where she has led the "Qualitative Diagnosis, Supervision and Control" Group for several years. Her main research interests are in Qualitative and Model-Based Reasoning and its application to dynamic systems Supervision and Diagnosis. She has been particularly active in bridging the AI and Control Engineering Model-Based Diagnosis communities, as leader of the BRIDGE Task Group of MONET. She has been responsible for several industrial and european projects and published more than 100 papers in international conference proceedings and scientific journals. Her current responsibilities include; member of the *IFAC SafeProcess* Technical Committee; member of the European Network of Excellence *MONET* Steering Committee; member of the French CNRS Network *RTP 20* on "Diagnosis, Reliability and Safety" Steering Committee. She is a Senior Member of the *IEEE* Computer Society.

Louise has been mountain climbing for 16 years, living in Toulouse, France, which is close to the Pyrénées Mountains. Technical rock climbing, technical ice climbing, mountain ski and mountaineering are among her activities. She has climbed a big wall in Yosemite and several rock routes in United-States, Sahara and Brasil. She participated in expeditions around the world, climbing 6000m peaks in the Andes; Ama Dablam (6812m) in the Himalayas, Nepal; and recently Aconcagua (6962m) which is the highest peak of the American continent. Closer to home, she has climbed numerous technical routes in the French and European mountains Pyrénées, Mont Blanc, Monte Rosa and Oberland. Louise has three children who share her rope on easy climbs.

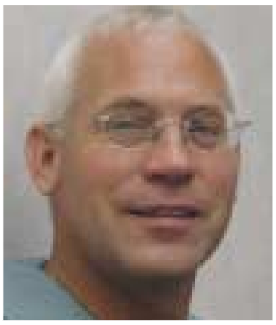

Dr. Robert Milne is the founder and Managing Director of Sermatech Intelligent Applications. Since the company was founded in 1986, it has focused on real time, on line diagnostic systems using Artificial Intelligence. The company has built many successful diagnostics systems and received a Queen's Award for Technology for their successful transition of research into industrial application. Currently the company's main product is Tiger, for real time gas turbine condition monitoring. The system has won many awards and is in use on 4 continents. Intelligent Applications was acquired by Sermatech Power Solutions in early 2001. Rob has also been a Director of the Scottish Software Federation for many years, is the President of ECCAI, the European Coordinating Committee for Artificial Intelligence, and is the Treasurer of the British Computer Society's specialist group on Artificial Intelligence. He has a BSc in computer science from the Massachusetts Institute of Technology and a PhD in Artificial Intelligence from the University of Edinburgh.

Rob has been mountain climbing for 30 years, having grown up in Colorado. His activities include technical rock climbing, technical ice climbing and mountaineering. He has climbed big walls in Yosemite, 14,000ft peaks in Colorado, 4000m peaks in the Alps (including Mt. Blanc, The Matterhorn and the north face of the Eiger). He has made first ascents of mountains in Alaska and Pakistan and with Aconcagua has been to the highest point on 4 of the continents. Others include Mount McKinley, North America; Kilimanjaro, Africa; and Carstensz Pyramid, Oceania. He has climbed all 284 peaks over 3,000ft in Scotland and recently co-authored the guidebook to "The Corbetts", peaks over 2,500ft in Scotland.

Editorial

It has been a busy six months for the Project and the Task Groups. The four Domain Areas have produced Technological Roadmaps and these have been incorporated into the Overall Roadmap for the Technologies, which is available online. This document will identify the likely industrial developments of the technologies, as well as the improvements to the technology that is needed to achieve the expected industrial development. Both will make a significant impact on the projects the Network is engaging in, see Task Group Reports.

Two members of the MONET Steering Committee are also editing a Special Issue of AI Magazine, for the AAAI organisation and this will contain articles covering all MONET areas. AI Magazine is a highly prestigious publication and has the widest publication coverage worldwide of any magazine in this field.

MONET has continued its clustering activities with EUNITE - see the Bio-Medical Task Group Section for details. Iain Russell gave a presentation on MONET to the General Assembly Meeting of EUNITE in July 2003 in order to foster closer links in the areas of Bio-Informatics and Bio-Medicine. The Bio-Medical Task Group organised a joint MONET - EUNITE Workshop on Bio-Informatics (Chaired by George Coghill) at the same event.

Here in the Project Office we have been streamlining and updating the MONET website and the MONET Information Resource. In addition to this, preparation for the forthcoming MONET Summer School in September is well under way, for details see the centre pages. On a similar note, MONET has arranged for a student, Jose Alberto Maestro Prieto, from the University of Valladolid to spend three months at the University of Wales, Aberystwyth under the direction of Professor Price.

Learning Qualitative Metabolic Models.

George M. Coghill * Simon M. Garrett Ross D. King †

Abstract

The ability to learn a model of a system from observations of the system and background knowledge is central to intelligence, and the automation of the process is a key research goal of Artificial Intelligence. We present a model-learning system, developed for application to scientific discovery problems, where the models are scientific hypotheses and the observations are experiments. The learning system, QOPH learns the *structural* relationships between the observed variables, known to be a hard problem. QOPH has been shown capable of learning models with hidden (unmeasured) variables, under different levels of noise, and from qualitative or quantitative input data.

1 Introduction

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 acheiving 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 paper we present QOPH 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.

*Department of Computing Science, University of Aberdeen, Aberdeen AB24 3UE. email: gcoghill@csd.abdn.ac.uk

†Department of Computer Science, University of Wales, Aberystwyth SY23 3DB. emai: {smg, rdk}@aber.ac.uk

2 Background

2.1 Qualitative Simulation

QSIM [8] 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.

In QSIM, each model consists of a set of variables linked together via a set of *constraints*, called a *qualitative differential equation* (QDE). Each variable consists of a $\langle qmag, qdir \rangle$ pair. Here, *qmag* is the qualitative magnitude of the variable. It has a quantity space of varying resolution consisting of alternating points (called landmark values) and intervals; typically the quantity space is divided into the regions $[-\infty \ldots 0)$, $[0]$, $(0 \ldots \infty]$, where infinity is treated as a value. A *qdir* is the qualitative rate of change of the variable, which has a fixed, three valued resolution (the three quantities being *inc*, for increasing; *dec*, for decreasing; and *std*, for steady). Each constraint has only one operation and is defined between two or three variables.

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.

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 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.

2.2 Inductive Logic Programming (ILP)

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:

$$B \not\models E \tag{1}$$

find a hypothesis H so that

$$B \wedge H \models E \tag{2}$$

Many solutions to this problem are possible, e.g. the trivial solutions of E, or B → E. The problem is therefore how to restrict solutions to suitable ones. In abduction [3] solutions are restricted to ground facts; in ILP more general solutions are allowed [11], 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 [10]; and we have previously applied machine learning and ILP to many scientific problems with success(e.g. [6]).

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.

2.3 Related Work

Automated model construction is an important and growing area of research which has as a central aim the provision of appropriate models for scientific and industrial tasks. The ideal situation would be for a learning system to be supplied with only positive data for some of the variables of the system of interest plus some background knowledge and then produce a model which explains the data in a physically meaningfully manner, identify any hidden (unmeasured) variables and not be over-constrained. This is the hard to achieve target at which researchers are aiming. Previous work in the area has tended to either require that all variables be measured (e.g. [5]), required negative data (e.g. [1]), generate models that were overconstrained (e.g. [7, 12]) or models that were logically but not physically equivalent to the plant being modelled (e.g. [1]). In addition there has been no comprehensive testing of the conditions under which learning of qualitative models is possible. For further details see Garrett *et al* [4].

3 Model Learning Methodology - The QOPH Method

The ALEPH ILP system [13] was used as a wrapper for the QOPH implementation, which was written separately. As with [1], we used a subset of QSIM, implemented in Prolog, as background knowledge for ILP. The task of the model learning method is to induce a model given example values for a known set of qualitative variables (a set of qualitative states), and the model language of qualitative relations that can be applied to those variables.

ILP, like much of learning, can be considered to be a search through a space of possible solutions. In the case of learning QSIM models, this space is the set of all possible QSIM models, partially ordered by generality. The relation-variable lattice is traversed by best-first search, and the search of this space can be constrained by the use of various heuristics. These heuristics can be generated from a number of sources: for example systems theory or the domain knowledge of the areas under investigation. In the former case the heuristics consist of general principles from systems theory such as: models must be parsimonious, operate under integral causality, and contain no algebraic loops (although these latter are preferences rather than absolute rules - since for some systems it is not always possible to achieve them). Also, for example, if one is working in the biological area some of the domain knowledge may consist of a set of rules regarding legal chemical reactions that may take place.

3.1 Testing the QOPH method

The QOPH system was developed as a tool to aid in the construction of structural models of systems in molecular biology. This is a domain in which data are sparse and inherently noisy; therefore it was important that QOPH be thoroughly tested under these conditions in order to ascertain its potential as such a tool; and the following set of experiments were devised for this purpose.

1. Starting with a complete envisionment (containing N states) every combination of $N - K$ states from the envisionment (for $K = 0 \ldots N$)

was created (giving an experiment space of $2^N - 1$ experiments) and the ability of QOPH to learn the target model from each set of states was tested. This set of experiments measures the sensitivity of QOPH to sparcity of data alone.

2. For the complete envisionment of N states, experiments were run in which the total number of states used to by QOPH to learn from was kept constant (at N) with the number of real states being progressively replaced by a number of qualitative noisy states; from 0 (no noise) to N (only noise). A noisy qualitative state is defined as a state that is not part of the complete envisionment but is of the same form, containing the same number and type of variables. This tests the supposed effect of noise introduced in the quantitative to qualitative conversion process.
3. For a selection of the experiments used in (1) a random number of qualitative noisy states were added to the real ones and the effect on learning measured. This was done to simulate the effect of converting noisy signals.
4. Finally experiments were run in which the whole process (from data acquisition and interpretation to model construction) for both clean and noisy data were performed.

In order to illustrate the approach used and the results obtained we will utilise a coupled two compartment model, since compartmental models are oftem used to represent metabolic systems. Details of the full set of tests and results can be found in [4].

In this system the input, $inflow_1$, is the input to compartment 1 and the output, $outflow_2$, is the elimination to the environemnt from compartment 2 (see Fig. 1).

The model of this system is:

$\mathrm{DERIV}(conc_1,\ netflow_1)$,
$\mathrm{DERIV}(conc_2,\ netflow_2)$,
$\mathrm{ADD}(conc_2,\ concDiff,\ conc_1)$,
$\mathrm{M}^+(concDiff,\ flow_{1-2})$,
$\mathrm{M}^+(conc_2,\ outflow_2)$,
$\mathrm{ADD}(netflow_2,\ outflow_2,\ flow_{1-2})$,
$\mathrm{ADD}(flow_{1-2},\ netflow_1,\ inflow_1)$.

Here there are three unmeasured variables: '$netflow_1$', '$netflow_2$', and '$concDiff$'. For the system to be correctly learned these variables will have to be induced. Variable '$inflow_1$' (the input) is exogenous to the model and so appears only once.

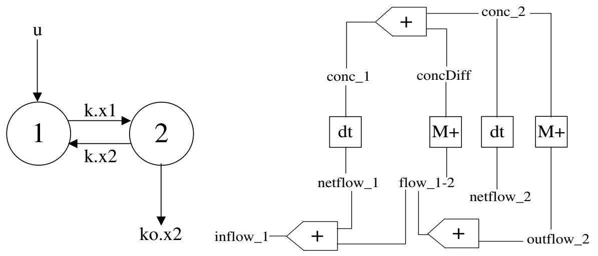

Figure 1: The coupled system (a) compartments; (b) QSIM

4 Results

Since any given experiment will induce its models from a finite number of states, it is possible to plot the *average* reliability for all the experiments for a particular number of states, from one state up to the number of states in the complete envisionment. This 'Average reliability' is given in the range [0 1]. For the noise experiments, the noise dimension is projected on the comparative 2-D plot (this assumes an average noise for each point on the state dimension) to allow comparison with clean data experiments, but a 3-D plot is also presented for the noise experiments that includes the noise dimension.

The plots of the number of states used against average reliability for the coupled tanks are shown in Fig. 2.

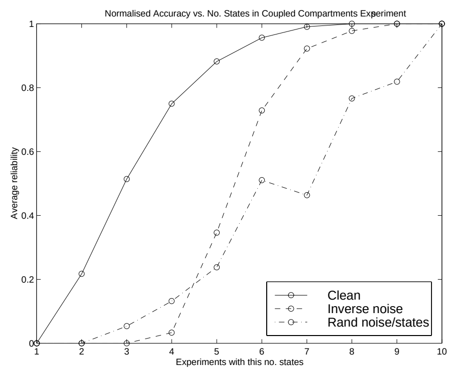

Figure 2: Coupled tanks reliability graphs

We analysed the performance of subsets of the complete envisionment to test whether certain subsets helped QOPH to learn the correct model more reliably than others. If this were the case then there would a a number of *minimal subsets* that contained the lowest number of states that reliably lead to the correct model being found. Subset analysis of the clean data experiments for the coupled tanks give the following states as the minimal subsets.

[1,6], [6,8], [6,9] (state 6 with 1, 8 or 9)
[2,8], ([6,8]), [7,8] (state 8 with 2, 6 or 7)
[1,2,3], [1,2,4], [1,2,5] (states 1 and 2 with 3, 4 or 5)
[1,3,7], [1,4,7], [1,5,7] (states 1 and 7 with 3, 4 or 5)
[3,7,9], [4,7,9], [5,7,9] (states 2 and 9 with 3, 4 or 5)
[2,3,9], [2,4,9], [2,5,9] (states 7 and 9 with 3, 4 or 5)

Fig. 3 shows the relationship of these states in the envisionment graph. A comparison with Table 1 reveals two key features: a selection of states from different behaviours and the use of the critical points of the system are the key to inducing the correct model reliably (see Discussion section below).

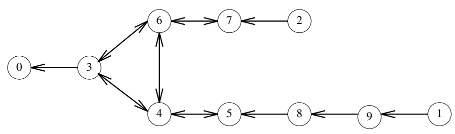

Figure 3: The envisionment graph for the coupled two compartment model

The results from the numerical data experiments are presented in Fig. 4. The legend in the top right corner associates initial values of the state variables (given as two concatenated digits) to a plot; 'all' is the case where the union of states from all initial conditions were used in learning. These results show that it is possible to learn models from clean and noisy numerical data. As discussed above, the qualitative states generated from the clean numerical data contain a number of unavoidable data transformation errors, and the resulting qualitative states form at most a single behaviour of the system under investigation. The set of states gleaned from quantitative to qualitative conversion did not form a full behaviour for the coupled two compartment model, which makes the ability to learn a model from them even more impressive.

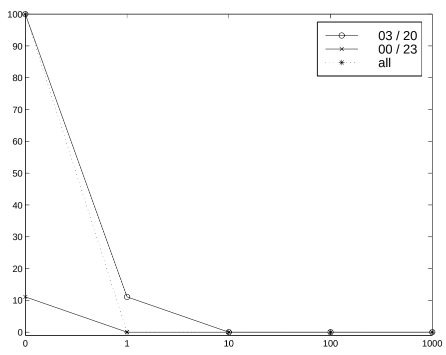

Figure 4: Reliability of learning the correct model from numerical data vs. 1000ths of full Gaussian noise for coupled tanks

5 Application to Biological Systems

As well as exploring the effects of sparcity of data and adding noise it was important to test the *scalablilty* of the QOPH learning method. So far we have only described models constructed from the basic QSIM primitives; to improve scalibility it was useful to be able to use the well-established AI principle of *chunking* [9].

Metabolic pathways essentially contain only two types of molecule: metabolites and enzymes, we therefore designed two *Metabolic Components*, built from standard QSIM relations, to model *metabolites* and *enzymes*. Concentrations of metabolites vary over time as they are synthesised or utilised by enzymatically catalysed reactions. This means that their concentration at time t is a function of their concentration at time $t-1$, and the amount that they are used or created by various enzyme reactions. This can be expressed as a simple summation in QSIM. The qualitative equation for the metabolite components is therefore:

$$\frac{dM}{dt} = M(t) + \sum_{i=0}^{n}(enzm\_flow_i). \qquad (3)$$

The other form of high-level metabolic component in a metabolic pathway are enzymes. Each enzyme is assumed to have one or two inputs and one or two outputs. If there are two inputs or outputs these are considered to form an input or output complex, such that the amount of the complex is proportional to the amount of the inputs or outputs multiplied together. The input complex is converted into the output complex which then disassociates into the output metabolites, and vice versa. The overall flow through the enzyme is the amount of input complex formed minus the amount of output complex formed. The qualitative equation for the enzyme components is therefore[1]:

$$flow = \underbrace{\mathbf{M}^{+}(\prod_{i=1}^{n} M_i)}_{inputcomplex} - \underbrace{\mathbf{M}^{+}(\prod_{j=1}^{m} M_j)}_{outputcomplex}. \qquad (4)$$

This is an abstraction of standard kinetic equations [2] and is an expression of the collision probabilities of the metabolites and enzyme. We assume for simplicity that enzymes are taken to exist in constant amounts; although this is clearly a simplification this assumption is also used in ODE modelling. These metabolic components are shown in Fig. 5.

A model of glycolysis in *Trypanasoma brucei* constructed from these Metabolic Components is

[1] Note the distinction between M and $\mathbf{M}^{+}$, the amount of a metabolite and the monotonically increasing relation respectively.

State	$level_A$	$level_B$	$crossflow_{AB}$	$outflow_B$
0	$< 0, std >$	$< 0, std >$	$< 0, std >$	$< 0, std >$
1	$< 0, inc >$	$< (0, \infty), dec >$	$< (-\infty, 0), inc >$	$< (0, \infty), dec >$
2	$< (0, \infty), dec >$	$< 0, inc >$	$< (0, \infty), dec >$	$< 0, inc >$
3	$< (0, \infty), dec >$	$< (0, \infty), dec >$	$< (0, \infty), dec >$	$< (0, \infty), dec >$
4	$< (0, \infty), dec >$	$< (0, \infty), dec >$	$< (0, \infty), std >$	$< (0, \infty), dec >$
5	$< (0, \infty), dec >$	$< (0, \infty), dec >$	$< (0, \infty), inc >$	$< (0, \infty), dec >$
6	$< (0, \infty), dec >$	$< (0, \infty), std >$	$< (0, \infty), dec >$	$< (0, \infty), std >$
7	$< (0, \infty), dec >$	$< (0, \infty), inc >$	$< (0, \infty), dec >$	$< (0, \infty), inc >$
8	$< (0, \infty), std >$	$< (0, \infty), dec >$	$< 0, inc >$	$< (0, \infty), dec >$
9	$< (0, \infty), inc >$	$< (0, \infty), dec >$	$< (-\infty, 0), inc >$	$< (0, \infty), dec >$

Table 1: The envisionment states for the coupled tanks.

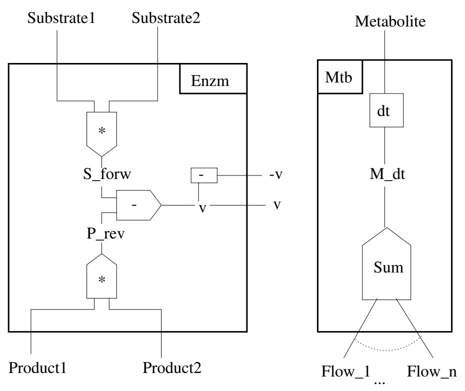

Figure 5: Metabolic components for metabolic system modelling

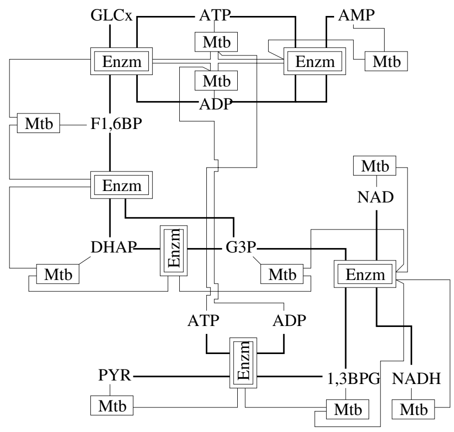

Figure 6: The glycolysis metabolic pathway, built from metabolic components

shown in Fig. 6. The qualitative model is easier to understand than an ODE since it extracts out detail and allows a complete envisionment of the states.

We gave QOPH the states for the top part of this system (shown below). There are 77 distinct legal qualitative states for this part of system.

Using just the states from the model and the metabolic components, described above, and assuming that no knowledge about chemical interactions between the molecules in glycolysis, it was possible to learn the top part of the model in only a few hours.

Learning a model of this sort represents a major step forward because the learned system is equivalent to a QSIM model consisting of 36 relations, It was calculated that introducing high level components has more than doubled the complexity of the models that can be learned, as well as making the resulting models easier to read.

6 Discussion and Conclusions

The first general point to note is that for all the experiments the number of measured variables from which learning took place remained constant and was less than the total number of variables in the target model. Thus in all circumstances the learning system had to find the hidden variables and their relationships to the other variables of the model.

Analysis of the clean data experiments showed that given the complete envisionment of a system the correct model was always reliably found. As one would expect there was a gradual deterioration in the reliability as the number of states presented as data was reduced. However, a closer analysis of the results in conjunction with the envisionment graphs for the target models reveals that there is a strong relationship between the reliability of the learning process and the number, and type, of states used in an experiment.

An interesting result from this analysis is the observation that models can be reliably learnt from a minimal number of qualitative states (two in the case of the coupled two compartment system) if the states come from different branches in the envisionment graph. So we can hypothesise first of all that in order to reliably learn a system the data used should come from experiments yielding qualitatively different behaviours (that is behaviours which would appear as distinct branches in an envisionment graph).

However, this hypothesis only provides a necessary, but not a sufficient condition for learning. It was noted during the analysis that in each case where the model was reliably learnt with a minimal number of states, at least one of the states is a critical point of the first derivative of at least one of the state variables: indicating the importance of these critical point states to the definition of a system. What this means is that if an experiment were set up in which all the state variables were exactly at their critical points then the experiment could be run for a very short time and the correct model structure identified. Of course, it is impossible to set up such an experiment, especially in the situation where the structure of the system is completely unknown. Another alternative is to set up multiple experiments with the state variables set to their extremæ: from these initial conditions all the states of the envisionment will eventually be passed through. The downside of this is that the experiments may be difficult to set up and could take an very long time to complete. These two scenarios form the ends of a spectrum within which the optimal experimental setting will lie; the identification of the the best strategies is an important area of research arising from the results of the present work, but it is beyond the scope of this paper.

The main results from this part of the project can be summarised as follows:

- The benchmark models could be induced from their complete envisionments.
- As the number of states chosen from the complete envisionment increases so does the frequency and reliability of finding the correct model.
- The correct model can always be reliably found given a relatively small subset of the total envisionment. There is a set of these subsets such that other state subsets are either supersets of one member of this set, or do not reliably give rise to the correct model.
- Even though subsets containing *very few* states can reliably give rise to the correct model, it is possible to select subsets containing *almost all* the states that do *not* reliably lead to the correct model.
- Qualitative models can be leart from data containing noisy qualitative states, though the overall reliability is reduced.
- Models can be learned from noisy simulated real data for the benchmark systems.

In addition to the results presented here we have also used QOPH to learn a qualitative model representing the complex biological process of glycolysis [4].

References

[1] I. Bratko and S. Muggleton. Learning qualitative models of dynamic systems. In S. Muggleton, editor, *Inductive Logic Programming*, pages 437–452. Academic Press, 1992.

[2] W. W. Cleland. The kinetics of enzyme-catalysed reactions with two or more substrates and products: I. nomenclature and rate equations. *Biochim. Biophys. Acta*, 67:104–137, 1963.

[3] P. A. Flach and A. C. Kakas. *Abduction and Induction: Essays on their relation and integration*. Kluwer Academic Publishers, 2000.

[4] S. M. Garrett, G. M. Coghill, R. D. King, and A. Srinivasan. On learning qualitative models of qualitative and real-valued data. Technical Report UWA-DCS-01-037, University of Wales, Aberystwyth, August 2001. Submitted to *Artificial Intelligence* Sept' 2001.

[5] D. T. Hau and E. W. Coiera. Learning qualitative models of dynamic systems. *Machine Learning*, 26:177–211, 1993.

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

[7] I. C. Kraan, B. L. Richards, and B. J. Kuipers. Automatic abduction of qualitative models. In *Proceedings of Qualitative Reasoning 1991 (QR'91)*, 1991.

[8] B. Kuipers. *Qualitative Reasoning*. MIT Press, 1994.

[9] J. E. Laird, P.S. Rosenbloom, and A. Newell. Chunking in soar: The anatomy of a general learning mechanism. *Machine Learning*, 1:11–46, 1986.

[10] S. Muggleton. *Inductive Logic Programming*. Academic Press, London, 1992.

[11] S. Muggleton and C. Feng. Efficient induction of logic programs. In *Proc. of the First Conf. on Algorithmic Learning Theory*. OHMSHA, Tokyo, 1990.

[12] A. C. C. Say and S. Kuru. Qualitative system indentification: deriving structure from behavior. *Artificial Intelligence*, 83:75–141, 1996.

[13] A. Srinivasan. *Aleph web site* :. web.comlab.ox.ac.uk/oucl/research/areas/machlearn/Aleph/aleph_toc.html, 2000.

Bio-Medical Task Group

During the last six months, Task Group members have been involved in approaching various commercial companies for information about the potential role model-based and qualitative reasoning approaches could play in their branch of industry. Contacts with research-oriented hospitals have also improved.

The Task Group has placed considerable effort into its Technological Roadmap. The Bio-Medical Domain is vast and varied with many factors relying on public / political will instead of Technological ability. Considerable time during the preparation of this document was spent on understanding the various constituencies that can be distinguished in the health-care and biomedical communities.

The Task Group has been involved in two Workshops in the first half of this year and has a further Workshop towards the end of this year.

EUNITE 2003 - Workshop on Intelligent Technologies for Gene Expression Based Individualised Medicine. 9th May 2003. Jena, Germany

EUNITE 2003 - Joint MONET - EUNITE Workshop on Bio-Informatics held at the European Symposium on Intelligent Technologies, Hybrid Systems and their implementation on Smart Adaptive Systems. 10th - 11th July 2003. Oulu, Finland

AIME'03 Workshop: 'Qualitative and Model-based Reasoning in Biomedicine' held as part of The European Conference on Artificial Intelligence in Medicine (AIME'03), 18th - 22nd 2003 October. Cyprus. See advert below.

AIME'03 Workshop: *Qualitative and Model-based Reasoning in Biomedicine*

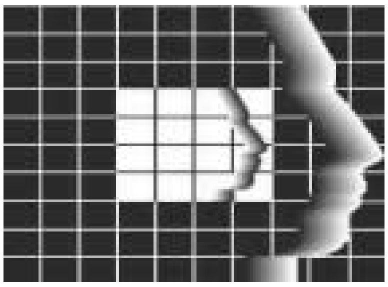

during ***The European Conference on Artificial Intelligence in Medicine*** (**AIME'03**)

Cyprus 19-22 October, 2003

(this CfP: http://www.cs.kun.nl/~peterl/aime03-ws.html)

- Submission of papers: **1 July, 2003**
- Notification of acceptance: 15 September, 2003
- Workshop: Sunday, 19th October 2003
- Conference: 19-22 October, 2003
- **With support from MONET**

To get further information about the AIME'03 Conference, please go to the following website: http://www.ucy.ac.cy/~aime03/

Automotive Task Group

This has been a very busy period for the Task Group and has seen them meet three times in the last six months. The majority of the work has been on investigating projects that could develop from the area and so far they have been involved in four such projects.

At an Automotive meeting in April the Task Group discussed the list of industrial requirements that had been gathered the previous year from several major European Automotive companies. The Task Group addressed the concerns from the Industrial partners and grouped these into four areas that could be interesting for future projects; they then gathered interest in these areas and worked towards generating projects. Four projects have emerged from this work concerning areas such as predictive diagnosis; test generation and Model-based testing of embedded systems; the application of intelligent systems / advance software to mobile systems; system autonomy (a joint project with the Space domain).

The Task Group have also been producing the first Version of their Domain specific Technological Roadmap in order to be able to offer continuing assistance to the competitiveness of the European Automotive Industry.

Also in June 2003 Louise Travé-Massuyès was invited to present at a Workshop on 'Immobot Technology and its Application to Automotive Realm' at the Car Internet Research Program (CIRP) Conference. CIRP is an industry-sponsored research initiative.

Model Based Fault Detection and Diagnosis (BRIDGE) Task Group

The BRIDGE Task Group has been working towards their Deliverables and Milestones and also putting considerable effort into reaching an ever increasing audience. They have also completed work on the Collection of Industrial Reference Problems which provides the description of a set of reference problems of different natures. These are presented along a common format making it easier to perform parallel analysis.

The focus of their Technological Roadmap (and also of the Task Group in General) is to provide both DX and FDI communities with a better understanding of the methodologies of the other community. This task will facilitate the future-planning that is shown in the Roadmap and the methods that are applied to achieving these goals will also be determined from the conclusions drawn in this document. The Task Group has also brought forward the work to produce tutorial materials as some have already been produced along with the Industrial Reference Problems and there will be more produced for the sessions on Diagnosis held during the MONET Summer School.

The Task Group held a one day BRIDGE workshop at the co-located International Workshop on Diagnosis 2003 and SafeProcess 2003. The day was attended by many members from the SafeProcess community and was a tremendous driver for increasing the awareness of researchers about FDI and DX technologies and their combination. Preparations are also beginning for the International Workshop on Diagnosis 2004 which will be held in Carcassonne (France) in 2004. Much of the organisation is being done by members of the BRIDGE Task Group.

The Task Group are also producing an IEEE BRIDGE Issue. The publication is now in the final stages of preparation and will contain six articles on BRIDGE issues and should appear around the second quarter of 2004.

Bert Bredeweg and Peter Struss (Guest editors)

AI MAGAZINE SPECIAL ISSUE

The qualitative challenge: Automated reasoning about the physical world

To appear: Winter 2003/2004

Contents

monet

MONET Summer School

The MONET Project is the European Network of Exellence in the Development of Model Based Systems and Qualitative Reasoning. Its three year life is dedicated to forwarding these Technologies and their transfer into Industry. MONET is holding a subsidised Summer School for Doctoral Students and Industrialist interested in increasing their knowledge of these Technologies. The Summer School will give an excellent foundation for research in the area or for applying these technologies to real world problems.

If you wish to participate in this Event, or have a student who would like to attend then please contact us.

Aldemar Cretan Village Hotel, Crete

6th - 12th September 2003

Programme Highlights

- Explanation of Technology by World Leading Experts
- Practical Technology Demonstrations
- Group Seminars
- Future Technology Planning
- Experts Available for Drop-in Clinics

For Details Please Contact:

Janet Thomas
monet@aber.ac.uk

Web: http://monet.aber.ac.uk

monet

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

Phone: 01970 628521
Fax: 01970 628536

MONET Summer School Venue

The hotel is a very well equipped and modern establishment, with all the facilities you would expect from such a venue. The restaurants had an abundance of choice with local specialities as well as themed evenings and the service is friendly and efficient. The main bar is outdoors, and undercover and there are also bars at each of the pools and a beach bar with an a la carte restaurant.

The conference facilities are on site and the main meeting room is large, well lit and has air conditioning. There is a large lobby area outside the conference room, which will be available for our exclusive use - a good place for coffee breaks and informal gatherings. There is also a terrace which runs along one side of the conference room.

Pool and beach facilities are excellent. There are four pools of varying sizes and locations within the complex, one of which has a waterslide. The beach area is clean and well kept, with sun-loungers and parasols which are free to hotel guests.

During the Event there will be an excursion to sample the Cretan way of life and experience the island's rich history and culture. This will culminate in a Gala Dinner at a local Taverna.

Weather in Crete is usually hot and sunny with temperatures of approx 25 degrees Celsius and water temperatures of 24 degrees Celsius during September.

Crete has a rich and varied culture. It was distinguished as the home of Europe’s earliest civilisation when excavations in the early twentieth century confirmed the legends of King Minos and the prehistoric Minoan society. In between times, Crete has had a diverse past, including a period of Turkish occupation until the Greek flag was raised over the island for the first time in 1913.

Summer School 2003

Timetable

<table>
<tr><th>Day</th><th>Session 1</th><th>Session 2</th><th>Session 3</th><th>Session 4</th><th>Evening</th></tr>
<tr><td>Saturday</td><td colspan="4">Delegates begin to arrive and settle into accommodation</td><td>Informal Gathering</td></tr>
<tr><td>Sunday</td><td colspan="3">Delegates continue to arrive and settle into accommodation</td><td>Registration</td><td>Welcome Event</td></tr>
<tr><td>Monday</td><td colspan="2">Plenary Lecture (Ken Forbus)</td><td colspan="2">Modelling Practical Session (Bert Bredeweg)</td><td>Drop in Session</td></tr>
<tr><td>Tuesday</td><td colspan="2">Ecology: Lecture & Practical Session (Bert Bredeweg)</td><td colspan="2">Diagnosis Lecture (Peter Struss / Louise Travé-Massuyès)</td><td>Drop in Session</td></tr>
<tr><td>Wednesday</td><td colspan="2">Diagnostic Practical (Peter Struss / Louise Travé-Massuyès)</td><td>Cognition Lecture (Ken Forbus & Dedre Gentner)</td><td colspan="2">Cultural Trip</td></tr>
<tr><td>Thursday</td><td>Education Lecture (Bert Bredeweg)</td><td>Bio-Medical Lecture (George Coghill)</td><td colspan="2">Industrial Application (Mark Lee / Neal Snooke)</td><td>Gala Dinner</td></tr>
<tr><td>Friday</td><td>‘Challenges in MBR’ Lecture (Chris Price)</td><td>Concluding Session – Panel with Ken, Bert, Peter, etc…</td><td colspan="3">Delegates depart</td></tr>
</table>

Institutions who have already registered Delegates

IST at Technische University of Graz
DAC TU Graz
Univerity of Nijmegen
University of Wales, Aberystwyth
Aberdeen University
Technische University of Berlin
Universitat Rovira i Virgili
Jena University
PIA, Kuwait
University of Valladolid
IRISA
University of Amsterdam
Brunel
Linköping University
Jaume I University
Technische University of Dresden

To enquire about attending the MONET Summer School please contact
Janet
At the MONET Project Office
on jnt @aber.ac.uk
or use the on-line application form

Names of Tutors

Bert Bredewg
George Coghill
Ken Forbus
Dedre Gentner
Mark Lee
Louise Trave-Massuyes
Chris Price
Neal Snooke
Peter Struss

QRSER: Qualitative Reasoning for Stream Ecosystem Restoration and Recovery Report from the First Workshop 6-8 March 2003 in Jena, Germany

Michael Neumann
University of Jena, Institute of Ecology, Dornburgerstr. 159, D-07745 Jena, Germany, m.neumann@uni-jena.de
and
Bert Bredeweg (Univ. of Amsterdam); **Paulo Salles** (Univ. of Brasilia); **Tim Nuttle** (Univ. of Jena)

Abstract
This report presents our initiative to start a pan-European collaboration to build qualitative reasoning models for stream ecosystem restoration and recovery. The first QRSER workshop was held in March 2003 in Jena, Germany, with 17 participants from Europe, Brazil and the USA. The program of the workshop was well balanced between theory and practice sessions. After an introduction to QR in general and to each participant's research, focused lectures about QR in ecology were given. Participants had the opportunity to work out assignments and to practice building QR models. A discussion about the future development of the QRSER initiative (see www.qrser.de), possible projects, proposals, and meetings created a successful end to the first QRSER workshop.

Introduction
Running water bodies constitute a major component of the landscape and are one of the most important limnological habitats with unique structure and function. They have, however, suffered decades of degradation from agricultural and industrial pollution and from weir or dam construction, channelisation, and dredging. This damage has reduced their ability to provide economically important ecosystem services including drinking water, irrigation, hydropower, commerce, recreation and fisheries.

Decision makers in restoration and recovery need predictions to plan their activities. Quantitative ecosystem models may provide such predictions, but they are not only difficult to parameterise, but also difficult to explain to non-experts. Stream ecosystems are especially difficult to model because of our generally poor understanding of ecosystem structure and functioning.

QR is an innovative technique that captures the fundamental aspects of a system, while sup-pressing much of the irrelevant detail. This approach makes expert knowledge available to non-experts for direct use in applied contexts. It will help reconcile the conflicting interests of water users and facilitate forecasting, management, and restoration of running waters throughout Europe.

Objectives
The QRSER workshop was held to bring together European stream ecologists and QR experts and to teach QR modelling techniques for ecological research. Consequently, the target audi-ence of this workshop were experienced stream ecologists from Europe who work in the area of regeneration, restoration, recovery, or improvement of degraded aquatic ecosystems. Ide-ally, they should work in a long-term research project and have already collected data and knowledge about the functioning of the investigated ecosystem.
Participants were invited to learn about the application of QR techniques to model ecological systems, share their own ideas about using QR models to investigate stream restoration and recovery, and practice thinking about representing familiar concepts in the intuitive, though unfamiliar, QR framework. Participants also practiced building QR models with example data for simplified systems. Another important goal of the workshop was to develop ideas for a proposal for the 6th framework programme of the European Union on collaborative research on QR modelling of stream ecosystems and to use QR to improve restoration and recovery in a variety of stream ecosystems.

The venue
Fifteen kilometres north of Jena, Germany, high above the valley of the river Saale, the three Dornburger Schlösser build a unique ensemble from different epochs and styles coming from the Middle Ages, Rococo, and Renaissance periods. They are connected together by vineyard terraces and romantic, rose-covered arcades. The workshop was held in the Renaissance Chateau, built ca. 1540. Goethe lived here in 1828 and so the building is named today "Goetheschloss", and represents one of the prettiest commemorative sites for this great poet.

The program
The program of the first QRSER workshop was divided into three parts. On the first day, a general introduction to QR was given. Participant had the chance to give a short intro-duction into their research projects and the area where they would like to build models with QR. It became clear that the focus of work encompassed diverse areas including water quality aspects, structure and size of communities, habitat quality, and integrated catchment man-agement.

During the second day of the workshop, an introduction about QR in ecology and practical assignments were given, followed by a more in-depth tutorial and more sophisticated assign-ment. A PC computer with the GARP software (including HOMER and VISIGARP) was available for every participant so that they could work on the assignments for a defined time on their own and the solution was presented to everybody at the

Prof Yordan Uzunov (Bulgaria): It was a perfect idea to hold this workshop introducing an emerging, innovative technique for non-numerical description of ecosystems of running water. The highest value of the workshop was the final commitment of the participants to join their expert knowledge implementing this. At the moment, the most advantageous feature of this approach is the opportunity to create many scenarios for possible developments within a stream ecosystem when introducing an integrated management and ecosystem approach.
Prof. Stefan Schmutz (Austria): State-of-the-art management of riverine ecosystems follows the principle of an integrated approach. Methods for integrating available knowledge into consistent frameworks are still widely lacking. Decision makers strongly need tools to integrate scattered information differing widely in expressiveness and quality. QR represents an adequate tool to handle these problems by developing consistent models for decision making.

end. This approach allowed each participant had the chance to learn more about the general background of QR in ecology and to practice building a QR model with GARP.

The third part of the workshop was a discussion about a proposal and the future devel-opment of the QRSER initiative. Participants discussed the idea, objectives and work plan of a project to be submitted to the European Commission 6th framework programme. All partici-pants of the workshop supported the idea of a specific targeted research project (STREP) proposal. A time schedule was developed for future collaboration and development of the QRSER initiative (see www.qrser.de). A second workshop at the end of 2003 and a third meeting in spring 2004 during an international conference was planned. The idea is to enable all participants to build QR models on their own and to present these models during an international meeting. Papers of this meeting will be reviewed and published in an edited book about qualitative reasoning in stream ecosystem restoration and recovery.

Overall, the program of the first QRSER workshop was well balanced between theory and practical classes to introduce a new and innovative technique to a group of stream ecologists from all over Europe. The program was flexible enough to accommodate the needs of the participants and left enough time for discussing and future planning. The organisers are very happy with the results of the first workshop and look forward to future developments.

Questionnaire to the participants

After the first QRSER workshop, 12 participants presented their comments about different aspects of the modelling effort by answering a questionnaire. Topics assessed were their individual knowledge and their opinions about the workshop contents, structure, and the model-building process. The overall evaluation of the QRSER structure and methodology was very good: "Friendly, relaxed atmosphere, good location, diverse array of participants willing to share ideas," "Group not too large", "Everyone had [the] opportunity to discuss, and over-domination of a few was kept minimal". The balance between theory and practice and the amount of material presented was considered about right, considering time available.

However, building QR models was not considered to be an easy task, even though most of the participants had 'some' previous experience modelling, and three of them were experts. Difficulties were found in different aspects of the modelling process. A new way of thinking about well known problems was required. For example, for half of those who answered the questionnaire, it was 'medium' or 'difficult' to create causal models of the system at hand, capture notions like direct and indirect influences imposed by processes, and to represent qualitatively objects, quantities and values. Technical aspects also presented some barriers: "difficult to learn the terminology and what [the modelling primitives] do'.

Asked about applications of QR to their work, the participants of the first QRSER workshop were enthusiastic: "Opportunities for immediate application in research activities" and "good application in management decision support in data limiting scenarios" were foreseen by the participants. Most of the answers mentioned educational applications of QR models: "very useful for education and describing general ecological behaviour common to all ecosystems". One of the answers may summarize the general feeling: "[it is very useful and] will influence my practical work".

All in all, we conclude that modelling using QR techniques is very much in line with ecological thinking. The participants indicated that they would recommend a QR approach to colleagues because they found it useful for research, management, and education in stream ecosystem recovery.

The QRSER Project for the EU FP6

Under the call FP6-2002-GLOBAL-1 with a deadline in April 2003, the consortium built during the first workshop submitted a proposal with the acronym "QRSER" for a STREP (reference number FP6-505353). The objective of the QRSER proposal is to apply QR to long-term investigations of stream ecosystems. Our consortium of ecologists and information technologists from 14 European countries and Brazil will develop a cost-effective technology to forecast stream parameters.

This is the first initiative to develop QR models of stream ecosystem restoration and recovery. These models will support researchers and managers throughout Europe to understand behaviour of these complex systems. Simultaneously, this project will develop ecological curricula to teach ecological concepts to university students and applied users, to increase ecological awareness and promote sustainable development. The work plan of this three-year project includes development of advanced software to support our collaborative approach. Ecological partners will develop QR models on specific problems, share them in a library of model fragments, and involve applied users for integrated model building.

The Relevance for Education

QR functions via the rich vocabulary that humans use to describe system behaviour, and thus facilitates causal reasoning. Hence, QR models are ideal for educational applications. Because of the various levels of understanding of learners, one model cannot be used at all levels of instruction. For example, learners who do not already have an advanced understanding of ecological theory cannot be expected to understand a complex community or ecosystem model if it is presented to them as a single entity. Rather, it is helpful to begin with simple, static models (descriptions of model components), and progressively add dynamic components (e.g., interactions between populations), until finally the full ecosystem model can be understood.

An important aspect of the QRSER initiative concerns its educational objectives. The idea is to develop university curricula using QR models to teach basic to advanced ecological concepts. Additionally, we have the goal to develop instructional methods to teach applied users (e.g. managers, decision makers and stakeholders) how to use QR models to understand and manage stream ecosystems to reach the goal of restoration and recovery.

Outlook for the Second QRSER Workshop

The second QRSER workshop is being organised by the Danube Delta National Institute for Research & Development, Tulcea, Romania. It is planned for the end of September 2003 (see www.qrser.de). The target audience are experienced stream ecologists from all over Europe. As a continuation of the first workshop, participants will learn and practise the GARP software to develop QR models. The second meeting aims to present the large range of stream ecosystem descriptions concerning the relationships among biotic and abiotic components. The goal is to enable participants to build QR models themselves.

An important objective is to involve end users. Participants from Water and Environment Protection Ministries, Water Agencies, etc. will learn and understand this new technique to represent knowledge about stream ecosystem behaviour. The goal is to support their decision making process for restoration and recovery. The EU Water Framework Directive and Flora Fauna Habitat Directive will be summarised and presented to the group to emphasise the current regulatory environment encompassing European stream-ecosystem quality.

The deadline for registration is 1 August 2003. For Information and instructions please contact Dr. Eugenia Cioaca from the Danube Delta National Institute for Research & Development (DDI); 165 Babadag Street, Tulcea, Romania (Tel.: +40240524550; E-mail: eugenia958@yahoo.co.uk).

Acknowledgements

The first QRSER workshop was financed by the Graduate Research Group: Analysis of the Functioning and Regeneration of Degraded Ecosystems (German Research Council project number GRK 266/1-96) at the University of Jena, Germany. In addition MONET provided travel support, via the educational task group, to facilitate the training during the practical classes.

Education and Training Task Group

The Education and Training Task Group has put considerable efforts, over this six month period, into starting their first project. They organised a Workshop in Jena, Germany and began work on a QR Stream Modelling system coupled with a Decision Support System for Ecosystem Management and Training. In other words this is effectively two projects rolled into one. The work that they have also done on the Roadmap and the Review of MBS&QR Technologies, along with the growth in the Task Group should see a great success in this and future projects.

Another major aspect of their work has been their Domain Technological Roadmap. Model-based Systems (MBS) and Qualitative Reasoning (QR) technology is of great importance for developing, strengthening and further improving education and training on topics dealing with systems and their behaviours. It is well known that an essential part of modern education and training involves the comprehension of systems and their behaviours. That is, being able to distinguish a system from the environment in which it operates, to identify the parts that it is made of, and to predict or explain its behaviours. The latter is concerning the overall system behaviour, how it potentially interacts with the behaviour of the environment, as well as how the system's behaviour originates from the individual behaviours of each of the parts that it consists of. The term 'Parts' does not only refer to components that constitute artefacts (e.g., a switch being part of a device of some kind), but also refers to entities or objects in general from which systems have grown (e.g., a tree being part of a forest) or have been built (e.g., a pole supporting some physical structure).

Although not all education and training involves reasoning about systems (for example, language teaching does not seem to fall within such a paradigm), many matters in educational settings, work situations and every day life do. MBS&QR technology is important because it provides a computer-based means to capture and communicate knowledge and insights that overcome limitations of currently used technology, such as numerical-based simulations. However, MBS&QR technology is not well known to a wider audience and there are few who are ready to use products and tools available to exploit the capabilities of this technology. This is an undesirable situation as a great possibility for improvement is potentially left unexploited. In this roadmap they depict the current situation concerning MBS&QR technology for education and training. They then envision how this should change in the future and how they may accomplish that vision; following the needs and desires from the field.

They have also been working to formulate a basis to practically demonstrate the applicability of MBS & QR technology to the Education and Training Domain. With the assistance of other Task Group Members, Bert Bredeweg conducted discussions with experts in the field of Ecosystem modelling and persuaded them that not only could MBS & QR build their models for them, but that it could also develop a 'Decision Support' system to assist in stream restoration and recovery programmes. For a full description of this (Jena) Project please see the article on pages 14 - 16.

The Task Group has also completed work on the (ED2) Review of MBS&QR Technologies and their application to teaching. Educational Model-based reasoning has a comparatively long history in model-based reasoning, starting with the Sophie system for teaching engineers to diagnose circuits in the early 1980's. This document examines the latest model-based technology being employed for tutoring and education-oriented systems, with case studies of some of the most recent systems to be built.

A Model-based Approach to Improved Prescription of Antibiotics

P.J.F. Lucas[1], C.A.M. Schurink[2], I.M. Hoepelman[2,3], M.J.M. Bonten[2,3]

[1]Institute for Computing and Information Sciences, University of Nijmegen
Nijmegen, The Netherlands, E-mail: peterl@cs.kun.nl
[2]Department of Medicine, Division of Acute Medicine and Infectious Diseases, and
[3]Division of Medical Microbiology, Eijkman-Winkler Institute
University Medical Centre, Utrecht, The Netherlands, E-mail: K.Schurink@azu.nl

Abstract

Bayesian networks have been introduced in the 1980s. Research to explore the use of the formalism in the context of medical decision making started in the 1990s. The formalism possesses the unique quality of being both a qualitative and quantitative, statistical knowledge-representation formalism. As it allows for structuring domain knowledge, by exploiting causal and other relationships between domain variables, the formalism is also model-based. In this paper, a Bayesian-network model of ventilator-associated pneumonia and an implementation of the decision-support system that incorporates this model and that is currently being evaluated in the ICU of the University Medical Centre Utrecht are described.

1 Introduction

The project described in this paper was initially undertaken to investigate the potential of the commercial clinical information system C2000 to act as a foundation for medical decision support in the ICU.[1] A 1994 study of antibiotics usage in Dutch ICUs revealed that 49% of the antibiotics were prescribed for respiratory-tract infections. As a clinical problem for the project the diagnosis and treatment of pneumonia in mechanically-ventilated patients was therefore chosen, which may be seen as an instance of a much wider clinical problem: the clinical management of infectious disease in hospitals. The significance of this derives from the presence of multi-resistant bacteria in clinical wards, in particular the ICU, makes prescription of antibiotics with a spectrum as narrow as possible essential; the prescription of broad-spectrum antibiotics promotes the development of antimicrobial resistance, and should therefore be avoided when possible. Most infectious-disease specialists and microbiologists therefore believe that the guidance of non-expert doctors in treating infectious disease must be improved; one way to achieve this aim may be through decision-support systems. A number of studies indicate that decision-support tools may indeed contribute to improving infectious-disease management and control [3, 5, 9, 15].

In our project *Bayesian networks* (BNs) have been chosen as the basis of most of the work. They have been introduced in the 1980s as a formalism for representing and reasoning with models of problems involving uncertainty, adopting probability theory as a basic framework [10]. Since the beginning of the 1990s researchers are exploring its possibilities for developing medical applications.

The BN formalism offers a natural way for representing the uncertainties involved in medicine when dealing with diagnosis, treatment selection, planning, and prediction of prognosis [6]. This is due to the fact that the probabilistic influences and interactions among variables can be described readily in a BN. As the formalism is declarative in nature, any (often conditional) probabilistic statement can be computed from a given BN, where the statement may concern both single and arbitrary Boolean combinations of variables. This allows asking questions such as "What is likely to be the result for the patient if I decide to request this test, or to prescribe this treatment". Another attractive feature of the formalism is that it is closely related to causal qualitative models, which explains why some researchers refer to it as the *causal probabilistic network (CPN) formalism*. An actual BN can often be understood in terms of cause-effect relationships reflected in its structure. Finally, there also exists a fully qualitative version of the Bayesian-network formalism, so-called qualitative probabilistic networks (QPNs) [11]. This, therefore, allows developers to choose for a fully qualitative modelling approach or for even further mixing qualitative and quantitative information.

In this paper, a BN model that was developed to assist clinicians in the diagnosis and selection of antibiotic treatment for patients with pneumonia in the ICU is described [8]. The model is part of a distributed decision-support system that allows clinicians to request advice concerning individual patients. This system is currently being evaluated within the ICU of the University Medical Centre Utrecht, and is also described in this paper.

2 Modelling

Developing a model of a realistic medical problem is usually far from easy, and using Bayesian networks for this purpose offers no exception in this respect. As is the case with other representation formalisms, there are particular guidelines which facili-

[1]C2000 is sold by the Eclipsys Corporation, `http://www.eclipsys.com`

Figure 1: Example of a mechanically ventilated patient in the ICU. Many of these patients develop pneumonia.

tate developing a BN [7]. We start by summarising some facts concerning the problem of ventilator-associated pneumonia in the ICU.

2.1 Ventilator-associated pneumonia

Many of the patients in the ICU are severely ill, which contributes to the likelihood that these patients get pneumonia. One explanation for this is that the functionality of the immune system in these patients is diminished. In addition, many patients admitted to an ICU need respiratory support by mechanical ventilation (See Figure 1). These, and a number of other factors, promote the development of bacterial pneumonia [1]. Pneumonia is a common disease in ICU patients; ventilator-associated pneumonia (VAP) may arise in patients who are mechanically ventilated. Because of the wide-spread dissemination of multi-resistant bacteria in hospitals and ICUs in particular, with which patients start to become colonised after one or two days, effective treatment of VAP is seen as an issue of major concern.

Unfortunately, already diagnosing the presence of VAP in patients is difficult, as many of the signs and symptoms that occur in VAP also occur in other disorders. For example, fever is a very common finding in patients in the ICU, and is typical for pneumonia, but it is more often associated with urinary tract infection. Hence, choosing the 'right' therapy, i.e. selecting antibiotics that are effective against the causative organisms, without causing major side effects and that are as much as possible directed to the causative organisms only, i.e. have a *narrow* antimicrobial spectrum, is even more difficult in the face of this uncertain diagnosis.

2.2 The Bayesian-network model

Figure 2 gives an overview of the structure of the BN model of VAP which we developed. The structure of a BN can be designed using knowledge of known causal dependences, influences or correlations. All or part of these may be derived from knowledge of domain experts, obtained from descriptions in literature, or extracted from data using structure-learning algorithms. Formally, a Bayesian network $\mathcal{B} = (G, \Pr)$ is a directed acyclic graph $G = (V(G), A(G))$ with set of vertices $V(G) = \{V_1, \ldots, V_n\}$, representing stochastic variables, and a set of arcs $A(G) \subseteq V(G) \times V(G)$, representing stochastic dependences and independences among the variables. On the set of stochastic variables a joint probability distribution $\Pr(V_1, \ldots, V_n)$ is defined that is factorised respecting the (in)dependences represented in the graph:

$$\Pr(V_1, \ldots, V_n) = \prod_{i=1}^{n} \Pr(V_i \mid \pi(V_i))$$

where $\pi(V_i)$ stands for the variables corresponding to the parents of vertex V_i. One of the attractive features of BNs is that it is possible to combine information from various sources, for example starting with defining a probability distribution from one source, and then refining it using data.

An important role in the model is played by the temporal process of *colonisation* of the airways by bacteria. The fact that this process is temporal, is expressed by the interaction between duration of stay (hospitalisation) and duration of mechanical ventilation: both are positively correlated to colonisation with particular bacteria. It is in principle also possible to represent temporal knowledge by means of temporal arcs between the same variables at different points in time, but reasoning with such a representation, which resembles a Markov process, may be very demanding computationally. Other arcs have a causal reading without a strong temporal connotation. For example, aspiration of stomach content is another factor positively correlated to colonisation with particular bacteria. When a patient gets colonised with a particular bacterium, there is a certain probability that pneumonia will develop. Therefore, an arc is drawn from 'colonisation' to 'pneumonia'. Duration of mechanical ventilation and the immunological status of a patient influence the probability that pneumonia will arise as well; therefore, an arc is drawn from 'immunological status' and 'mechanical ventilation' to 'pneumonia'. When a patient is affected by pneumonia, symptoms and signs can be observed, as well as abnormalities in laboratory values; this part of the model is shown in Figure 3. Here the arcs sometimes have a causal reading and sometimes the less specific meaning of a correlational influence.

Graphs like the one shown in Figure 2 appear easy to understand, but their underlying formal semantics is sophisticated. For example, the structure in Figure 3 tells us that leucocytosis is conditionally independent of body temperature given presence or absence of pneumonia. The notion of *induced* dependence is also central to the theory; it signifies a (dynamic) change in the dependence relation represented by the graph. For example, the various colonisation variables are (unconditionally) independent, but will become dependent once infor-

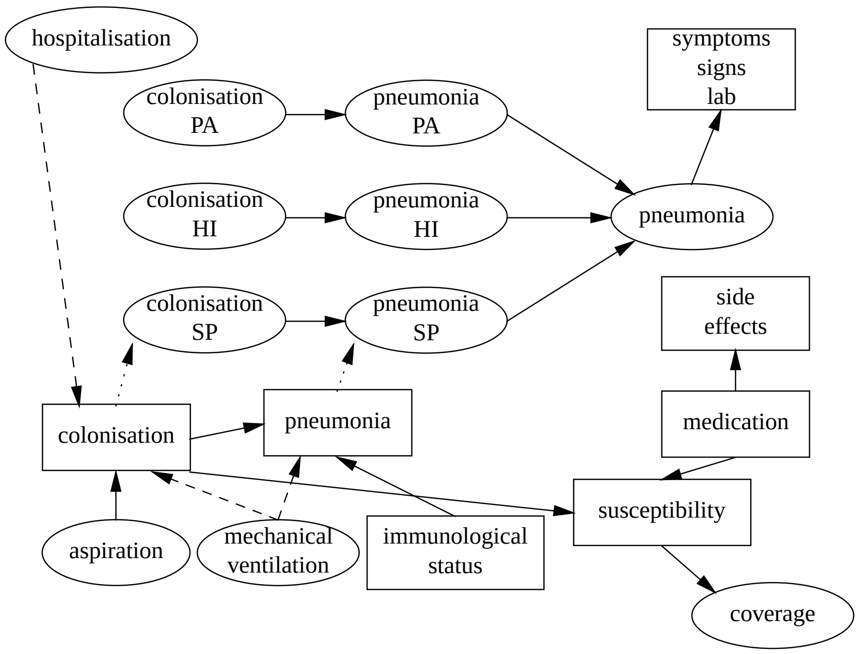

Figure 2: Detailed structure of part of the VAP model. Only three of the microorganisms included in the model are shown. Boxes stand for collections of similar vertices. Dotted arcs point to the actual topology of the network. Solid arcs stand for atemporal stochastic influences, whereas dashed arcs indicate temporal influences. Abbreviations of names of bacteria: PA = *Pseudomonas aeruginosa*, HI = *Haemophilus influenzae*, SP = *Streptococcus pneumoniae*.

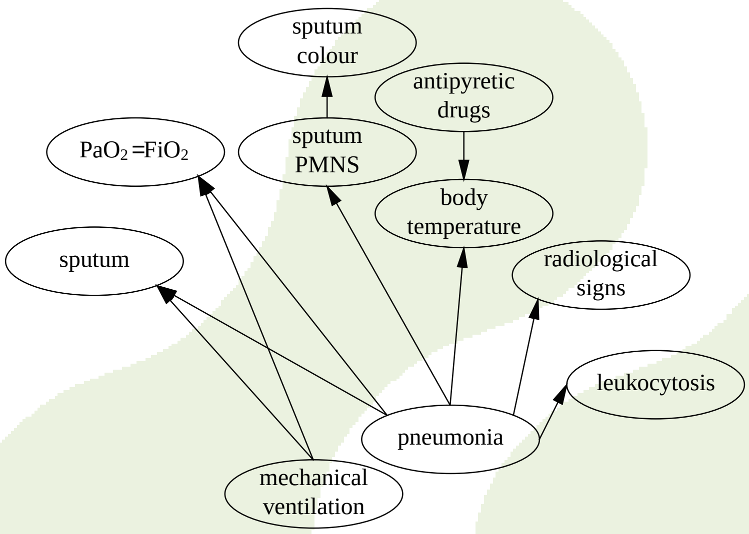

Figure 3: Probabilistic model of signs and symptoms of pneumonia.

mation on the presence or absence of their common consequence, pneumonia, is entered into the network. Insufficient understanding of the formal meaning of BNs may give rise to modelling flaws.

Prescription of antibiotic therapy amounts to selecting none, one or two antibiotic drugs, which was originally modelled by two identical therapy vertices. Let d be the number of possible drugs (including none). Then with two treatment vertices d^2 combinations are possible, of which $\binom{d}{2}^{1}$ are unique; the total number of different (bi- and mono-)therapies that actually can be prescribed is thus $\binom{d}{2}^{1} + d$. However, as not every possible combination makes clinical sense, a single treatment vertex was used in the final version of the network to represent the prescription of one and two antibiotics.

3 Medical problem solving

As a Bayesian network allows for the computation of any probabilistic statement, if all variables relevant for making a diagnosis and for prediction and treatment selection are included, the same network can be used to deal with a variety of medical-decision making tasks. This is an example of knowledge reuse; it will be illustrated below for the VAP model.

3.1 Diagnosis of pneumonia

Diagnosing VAP is a difficult task, because none of the signs and symptoms are unique for the bacteria that cause VAP. Determining a diagnosis based on available evidence $\mathcal{E}$ is often defined as:

$$d = \arg\max_{d \in D} \Pr(d \mid \mathcal{E})$$

where D here stands for the 'pneumonia' variable, and $\mathcal{E}$ for evidence, such as presence of leucocytosis, body temperature, duration of hospitalisation, and mechanical ventilation. Receiver operating characteristics (ROC) analysis is another frequently used method. It is employed to determine a probability cut-off point, which is then used to establish a diagnosis for future cases [14]. ROC analysis, however, requires a gold standard diagnosis, which often is not available in medicine. This is actually a problem with the diagnosis of VAP, as its pathological diagnosis is very unreliable. The results of an ROC analysis of the model with an infectious disease specialist and the ICU clinicians as gold standards are shown in Figure 4.

As mentioned above, the BN model of VAP incorporates temporal knowledge; however, it was recently shown that this is not really important for the diagnosis of VAP [2]. This can be understood by the fact that progress in time increases the likelihood of pneumonia, but time does not interact in a complicated non-monotonic fashion with 'pneumonia'. This implies that for the purpose of diagnosis it would be sufficient to use the part of the model shown in Figure 3, with a prior probability distribution for the variable 'pneumonia' determined by the marginal probability distribution as derived from the complete model.

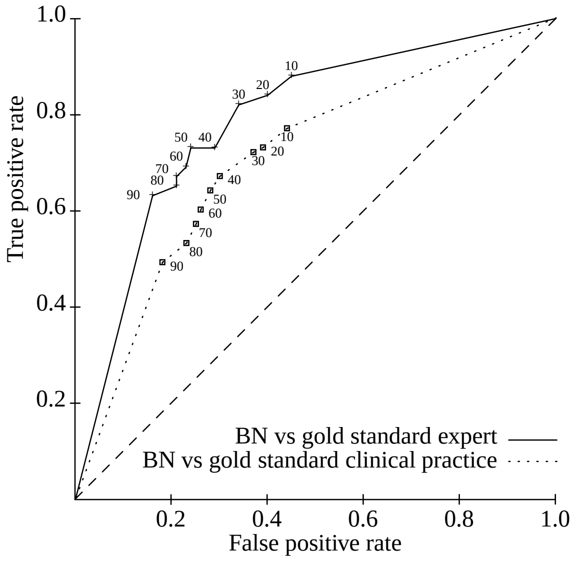

Figure 4: ROC curves based on patient data. The points are labelled with cut-off points in percentages. The upper curve is based on expert judgement; the lower curve on judgements by ICU doctors. The BN's performance is clearly nearer to expert opinion than to clinical practice.

3.2 Prediction and treatment selection

For the purpose of prediction of likely causative organisms, as well as for the selection of optimal antibiotic therapy, the temporal knowledge incorporated into the Bayesian-network model of VAP is of major importance. Figure 5 clearly indicates that both likelihood of colonisation and pneumonia by particular pathogens vary in time. As in particular the 'colonisation' variables together with selected antibiotics determine choice of treatment by predicting coverage, time cannot be ignored [2].

Treatment selection is based on selecting the antibiotic combination that yields an optimal outcome. In the case of treatment of VAP this can be defined as maximal coverage with minimal side effects, using antibiotics with a spectrum as narrow as possible, as this reduces the chances of the development of antimicrobial resistance in the hospital. This implies that the Bayesian network needs to be extended with decision theory, i.e. a *utility function*

$$u : \text{COVERAGE} \times \text{SIDE-EFFECTS} \times \text{SPECTRUM} \to \mathbb{R}$$

has to defined and treatment variables become *decision variables*. The resulting formalism is known under various names, among others *decision networks* and *influence diagrams* [12, 13]. The optimal treatment is the one with maximum expected utility.

Influence diagrams can be converted to Bayesian networks, among others by mapping the (bounded) image of the utility function u to the interval $[0, 1]$, and Bayesian-network inference algorithms can be used to determine (the sequence of) optimal decisions. In the VAP model, this mapping is very straightforward, as there is only one decision to make (antibiotic therapy). The actual mapping is derived in Ref. [8].

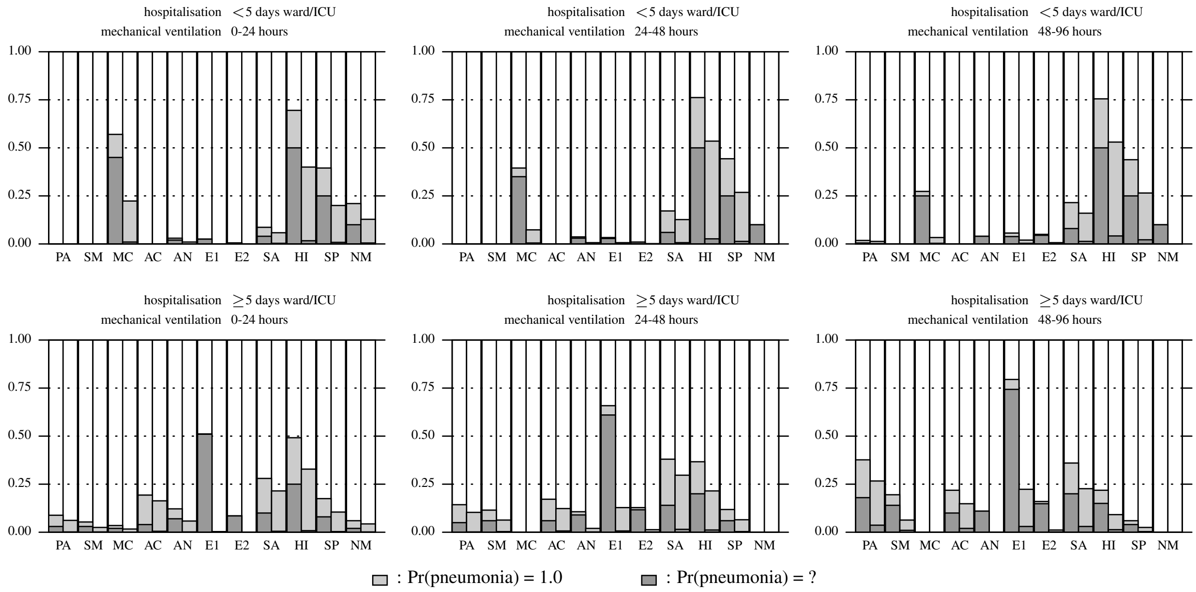

Figure 5: Obtained predictions after entering information concerning duration of hospitalisation and mechanical ventilation. Names of pathogens have been abbreviated (e.g. PA stands for *Pseudomonas aeruginosa* and SA for *Staphylococcus aureus*). For each pathogen, the probability of colonisation and pneumonia are depicted, in that order.

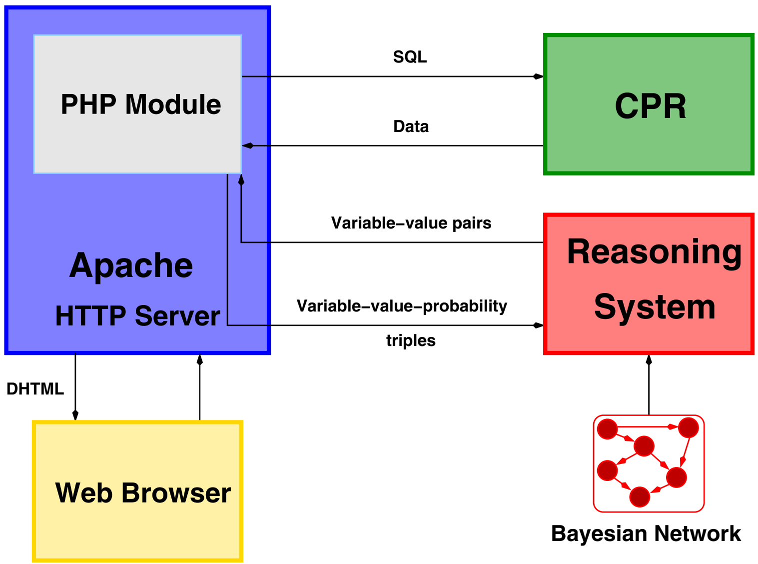

Figure 6: Architecture of the decision-support system that has been integrated with the clinical information system C2000 of Eclipsys. CPR stands for Computer-based Patient Record system, i.e. C2000.

4 Implementation and evaluation

A preliminary laboratory evaluation of the Bayesian network and decision-theoretic model has been carried out, and yielded promising results [8]. However, one of the major problems in the project has been the fact that VAP is not commonly recorded in C2000 by the ICU doctors, as VAP is never the reason for admission to the ICU but a concomitant disease in mechanically ventilated patients. In addition, there is not a single reliable gold standard for the diagnosis of VAP, and so the only way to make progress was to have each patient being judged on having VAP or not by one of the infectious disease experts. This has been taken into account in the design and implementation of the decision-support system, which has been set up in such a fashion that it supports carrying out clinical trials.

The overall architecture of the present decision-support system is shown in Figure 6. The system runs on a RedHat Linux server, which ensures that the decision-support system does not place extra CPU and memory load on the C2000 clinical information system servers. Information from C2000 to the decision-support system is extracted from the Sybase back-end of C2000 by SQL scripts. A PHP module takes care of the communication between C2000, Web clients (e.g. a Web browser used by the doctors), and the Bayesian-network reasoning engine. Hence, the decision-support system is accessible at every bed workstation from the C2000 graphical user-interface, and also from the hospital's intranet by those granted access to it.

Currently, the system is undergoing a clinical trial. The setup of the study is as follows. Before entering any information, the ICU doctor has to enter a clinical diagnosis and preferred antimicrobial treatment. Subsequently, the doctor has to enter part of patient-specific information; most of the information, however, is extracted from the C2000 patient records, and is simply presented to the doctor. On the average in 50% of the cases, the doctor is given an advice concerning diagnosis and treatment of the patient; in the remainder 50% no advice is given. The doctor is finally requested to enter preferred diagnosis and treatment again, and arguments for changes from the first entry. This set-up ensures that it is possible to filter out the *Hawthorne effect*, an effect on study outcome caused by the circumstance that the medical doctors know that their performance is being measured [4].

5 Conclusion

We have attempted to convey an impression of the process underlying the development and clinical deployment of a model-based decision-support system that intends to assist medical doctors in diagnosing VAP and selecting appropriate antimicrobial treatment for this disorder. The main advantages of adopting a model-based approach from a biomedical point of view are its versatility and strong links with how biomedical people think about problems. Also when data are not available, or scarce, as was the case in the early phases of our project, it is still possible to design a Bayesian network using subjective estimates based on expert knowledge. The subjective estimates can then be refined later when data become available.

Providing access to the decision-support system from every ICU-bed's workstation was an essential prerequisite for the success of our project, as was its integration with the clinical information system C2000. In the clinical trail we plan to study the effects of the system on the diagnostic and prescription performance of ICU doctors.

References

[1] J.G. Bartlett. *Management of Respiratory Tract Infections*. Williams & Wilkins, Baltimore, 1997.

[2] N.C. de Bruijn, P.J.F. Lucas, K. Schurink, M. Bonten, A. Hoepelman. Temporal probabilistic knowledge for medical decision making, in: *AIME 2001*. Lecture Notes in Artificial Intelligence 2101, Springer-Verlag, Berlin, 2001, pp. 231–234.

[3] R.S. Evans, S.L. Pestotnik, D.C. Classen, T.P. Clemmer, L.K. Weaver, J.F. Orme, J.F. Lloyd, J.P. Burke. A computer-assisted management program for antibiotics and other antiinfective agents. *N Engl J Med* 338(4) (1998) 232–238.

[4] C.P. Friedman and J.C. Wyatt. *Evaluation Methods in Medical Informatics*. Springer-Verlag, New York, 1997.

[5] L. Leibovici, V. Geitelman, Y. Yehezkelli, O. Poznanski, G. Milo, M. Paul, P. Ein-Dor. Improving empirical antibiotic treatment: prospective nonintervention testing of a decision support system. *J Intern Med* 242 (1997) 395–400.

[6] P.J.F. Lucas, H. Boot and B.G. Taal. Decision-theoretic network approach to treatment management and prognosis. *Knowledge-based Systems* 11 (1998) 321–330.

[7] P.J.F. Lucas, H. Boot and B.G. Taal, Computer-based decision-support in the management of primary gastric non-Hodgkin lymphoma, *Methods of Information in Medicine* 37 (1998) 206–219.

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

[9] W.J. McIsaac and V. Goel. Effect of an explicit decision-support tool on decisions to prescribe antibiotics for sore throat. *Med Decis Making* 18(2) (1998) 220–228.

[10] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman, San Mateo, California, 1988.

[11] S. Renooij. *Qualitative Approaches to Quantifying Probabilistic Networks*. PhD Thesis, Utrecht University, 2001.

[12] S.D. Shachter. Evaluating influence diagrams. *Operation Research* 34(6) (1986) 871–882.

[13] J.Q. Smith. *Decision Analysis*. Chapman and Hall, London, 1988.

[14] H.C. Sox, M.A. Blatt, M.C. Higgins, K.I. Marton. *Medical Decision Making*. Butterswort h, Boston, 1988.

[15] H. Warner Jr, S.R. Blue, D. Sorenson, L. Reimer, L. Li, M. Nelson, M. Barton, H. Warner. New computer-based tools for empiric antibiotic decision support. *JAMIA* (1997) 238–242.

The17th International Workshop on
Qualitative Reasoning
20-22 August 2003 in Brasília, Brazil

The Programme will include sessions on the following:

- "Fundamentos de Raciocínio Qualitativo e de Raciocínio Baseado em Modelos"
- "An introduction to ecological modelling using the Simile visual modelling environment"
- "Aplicações de QR/MBR à identificação de sistemas e à análise de circuitos elétricos"

- Workshop on QR and Financial Analysis (QR-$)
- Workshop on QR and Education (QR-Ed)
- 17th International Workshop on QR (17th QR)

- Symposium QR and Water Resources Management (QR-H2O)

For more information and application details, please look at the following website:

http://www.unb.br/ib/necbio/QR03/

MONET Information Resource (MIR)

MIR is a database in which MBS and QR reference are stored. The information is in four categories; Authors, Publications (search by title or abstract), Lecture Notes, and Applications. Access to the database is free to anyone who is associated with a MONET Member Institution.

Reference Count			
Authors	Publications	Lecture Notes	Applications
1237	1433	25	4
–	78 Full pdf's	23 Full pdf's	–

Resource Submission

We are continuously updating MIR with references, but if you have any papers, lecture notes or applications that you would allow us to put on the database please either email them to us or use the on-line submission form on the MONET Webpage.

MONET Project Office

Department of Computer Science
The University of Wales, Aberystwyth
Aberystwyth
Ceredigion
Wales SY23 3DB

Tel: +44 (0)1970 628521
Fax: +44 (0)1970 628536
Email: monet@aber.ac.uk
Webpage: http://monet.aber.ac.uk

monet