

Biomedical Domain Status Document

MONET 2 (IST-2001-33540), Biomedical Task Group

Deliverable B1 (version 2)

Contents

Contributors	ii
Purpose of this document	1
1 Introduction	2
2 Model-based approaches in biomedicine	2
2.1 Diagnostic systems in medicine	3
2.1.1 The problem of diagnosis and early approaches	3
2.1.2 Formal theories of diagnosis	3
2.1.3 The nature of medical domain models	8
2.2 Prognostic systems in medicine	9
2.2.1 Prognosis and outcome	9
2.2.2 Representation and construction of prognostic models	9
2.2.3 Traditional approaches: decision analysis and statistics	9
2.2.4 Model-based methods	12
2.2.5 Clinical usefulness	14
2.3 Treatment selection in medicine	15
2.4 Qualitative modelling and simulation in bioinformatics	16
2.4.1 Challenges to molecular biology	16
2.4.2 Qualitative reasoning and simulation	16
2.4.3 Bayesian networks and the discovery of functional interactions	17
2.4.4 Upscalability	17
3 Model-based and qualitative reasoning methods in biomedical industry and health-care	18
3.1 Reprogramming pacemakers	19
3.2 Challenging QR research issues in electrocardiology	21
3.3 A model-based approach for learning to identify cardiac arrhythmias	23
3.4 Qualitative reasoning for the modelling of metabolic systems	26
3.5 Advice in the prescription of antibiotics	27
4 Conclusions	28

Contributors

George Coghill
Raffaella Guglielmann
Liliana Ironi
Hidde de Jong
Ruud Kuipers
Peter Lucas
Marie Odile-Cordier
Stefania Tentoni
René Quiniou

Purpose of this document

This document is the result of work carried out by the Biomedical Task Group of the MONET 2 project (IST-2001-33540) in gathering information about the status of model-based and qualitative reasoning technology in the biomedical domain.

This document provides:

- a rationale of why model-based and qualitative reasoning technology are worthwhile exploring in the context of biomedicine;
- an overview of various successful model-based and qualitative-reasoning research results in biomedicine;
- an impression of the status of the use of model-based and qualitative-reasoning technology in health-care and biomedical industry.

This document is aimed at fulfilling the requirements for deliverable B1 (version 2) of the Biomedical Task Group of MONET 2.

1 Introduction

Modelling has always been one of the central approaches when tackling problems in medicine and biology, as in these fields it often pays off to take into account explicit knowledge of structure and underlying mechanisms. In many cases, there is a basic, *qualitative* understanding of biological mechanisms underlying processes; the high complexity and the non-standard mathematical nature of biological mechanisms, however, makes it hard, and often even impossible, to describe these mechanisms using quantitative techniques, except roughly in terms of statistics. As the understanding of the nature of biological processes is fundamental to life sciences, it is not surprising that much work in these fields is essentially *model-based*. However, the biomedical community is mostly not familiar with what qualitative and model-based reasoning (QR & MBR) technology has to offer. Consequently, there is a clear role for the Biomedical Task Group of MONET 2 to increase the awareness of biomedical researchers, industry and health-care organisations of the potential of model-based and qualitative-reasoning technology.

In this report, we sketch the historic development of the use of model-based and qualitative-reasoning methods in biomedicine, and look upon this technology from the viewpoint of modern biomedical research. Biomedical research is currently very much driven by empirical evidence and the analysis of experimental data, which is why we have attempted to take this emerging setting of biomedical research and health-care organisations into account. Until recently, clinical patient data have been stored in almost inaccessible paper form, hampering scientific use of this information. Furthermore, the IT infrastructure of hospitals consisted for more than 25 years only of a hospital-wide information system, storing administrative data and laboratory data. However, this infrastructure is currently being revolutionised by the introduction of electronic communication and clinical information systems in the wards. Similar developments are taking place in primary care. Also bioinformatics is giving a major impetus to fundamental research in the Life Sciences, where again being able to interpret data is of primary importance. These developments may offer new opportunities for qualitative-reasoning and model-based techniques, as these techniques are very suitable for bridging the gap between raw data, or the lack thereof, and constructing understandable models that can be manipulated to solve problems.

The report is structured as follows. In Section 2 we summarise research in using model-based reasoning methods in biomedicine. Section 3 summarises our findings on the use of qualitative-reasoning and model-based techniques in biomedical industry. Finally, in Section 4 we comment on our findings.

2 Model-based approaches in biomedicine

We start by reviewing model-based diagnosis in medicine, as model-based diagnosis is a particularly well-developed technology, and many medical applications of it have been and are being developed. Subsequently, the use of model-based and qualitative approaches in predicting outcome of actions is considered. Finally, the exploitation of model-based and qualitative reasoning in bioinformatics is reviewed.

2.1 Diagnostic systems in medicine

2.1.1 The problem of diagnosis and early approaches

Discovering what is wrong in a patient with specific symptoms and signs, i.e. diagnosis, is usually the starting point of medical management. It is, therefore, not surprising that diagnosis was one of the first subjects studied in the field of Medical Artificial Intelligence: in the burgeoning field of expert systems in the 1970s and 1980s, medical diagnostic applications abound. In fact, the impact of systems like MYCIN, [109], and INTERNIST-1, [87], on the field of Artificial Intelligence in general has been so large that many people outside the medical field believe that diagnostic expert systems are now in routine use in almost every hospital.

Only few systems developed in the early years have actually reached the stage of clinical maturity. Examples are a system for the diagnosis of acute abdominal pain [27], PUFF, a system for the automated interpretation of pulmonary function tests [4], and HEPAXPERT-I, a system that is used for the interpretation of serological tests for viral hepatitis A and B [3]. There have been various reasons for the limited spread of medical diagnostic systems. Firstly, the need of clinicians for computer-based decision support in the diagnostic process appears less prominent than thought by many researchers in the early days of Medical Artificial Intelligence. Secondly, the scientific basis of diagnosis was in a premature stage in the 1970s and 1980s. In critical fields like medicine, techniques lacking a solid foundation are not readily accepted. Thirdly, developing a reliable diagnostic system is a much demanding job, often requiring more time, effort and money than both researchers and funding agencies are prepared to invest. Obviously, this is still a problem today. Finally, the computational infrastructure in hospitals is often inadequate, and most physicians are not even accustomed to the idea of computer-aided problem solving.

Currently, research into medical diagnostic systems is again gaining momentum, in particular in well-chosen medical niches where the diagnostic process can be improved significantly by using reliable diagnostic systems. For example, the early clinical diagnosis of jaundice or ventilator-associated pneumonia are widely recognised as being difficult tasks for non-expert clinicians, requiring some form of decision support [118, 82]. Much progress has been made in the last decade in the development of solid theoretical foundations for diagnostic systems, employing logic, probability theory and set theory, thus providing a richer and more mature starting point for developers. Improvements to the theory of diagnostic systems are being developed up to this day. An important shift in the research of the last decade has also been that modern medical diagnostic systems increasingly rely on models of disease processes and models of structure and function of the human body. Furthermore, the infrastructure of health-care computing is currently being improved. As a consequence of the factors above, the outlook for diagnostic systems that are adequately integrated in the clinical environment is now much better than a decade ago [110], and this is reflected by the fact that many researchers are again involved in the development of such systems. People involved in such research now also see possibilities for developing such research results into products.

2.1.2 Formal theories of diagnosis

Many early diagnostic systems in medicine were based on empirical knowledge, often expressed as rules, and used in conjunction with a deductive reasoning method to solve diagnostic problems. Such rules were interpreted as being part of a classification relation; patients with specific symptoms and signs could be classified as belonging to certain diagnostic categories in terms of this relation [80].

The insight that pathophysiological models could also be employed as a basis for diagnosis already emerged at the end of the 1970s in the CASNET project [127]. In the CASNET system, medical




Figure 1: Basic idea underlying consistency-based diagnosis.

knowledge was expressed as (atemporal) cause-effect relationships, modelling the natural course of disease. ABEL was another early system incorporating causal knowledge; it was capable of diagnosing acid-base and electrolyte disturbances in a patient [93]. The major lesson learnt from these projects is that the representation of domain knowledge and the interpretation of this knowledge for a particular task, such as diagnosis, can be separated; hence, diagnosis need not necessarily be done in terms of a straightforward classification relation.

At the beginning of the 1980s, J.A. Reggia and colleagues proposed a new formal theory of diagnosis based on set theory, called the *set-covering theory* [104]. This theory offers a precise, formal definition of diagnosis of a problem using causal knowledge concerning disorders. Subsequent work has yielded several algorithms to compute set-covering diagnoses efficiently [96, 107, 129].

R. Reiter proposed a *logic-based* formal theory of diagnosis in 1987, aiming at formally capturing diagnosis of abnormal behaviour in a device or system, using a model of normal structure and behaviour [105]. Nowadays, Reiter's theory, which was later extended by J. de Kleer and colleagues, [65], to deal with knowledge of both normal and abnormal behaviour, is usually referred to as the theory of *consistency-based diagnosis* (See Figure 1). Basically, consistency-based diagnosis amounts to finding faulty device components that account for a discrepancy between predicted normal behaviour of a device, according to a model, and actually observed behaviour. The discrepancy is formalised as logical inconsistency; a diagnosis is established when assuming particular components to be faulty and others to be normally functioning restores consistency.

Approximately at the time of the publication of Reiter's seminal article, another logical definition of diagnosis was proposed by several researchers [25, 97]. This time the underlying knowledge used for diagnosis was assumed to consist of causal relationships. The theory is known as the *theory of abductive diagnosis* (See Figure 2), because diagnosis is formalised as reasoning from effects to causes, using logical implications of the form

$$\text{causes} \rightarrow \text{effects}$$

This type of reasoning is referred to as *abduction*; it is contrasted with deduction, which for implications of the form above, would amount to reasoning from causes to effects (provided that no negations of *effects* are given). Although the set-covering theory mentioned above is expressed in set theory and abductive diagnosis uses logic, the two theories have much in common. The theory of abductive diagnosis, however, is more expressive, because it is possible to explicitly represent various sorts of interaction among disorders, which is not possible in the original set-covering theory. For example, it is not possible to express in the original set-covering theory that the simultaneous occurrence of two or more disorders leads to masking of certain signs. However, these limitations of set-covering theory can be easily redressed, as shown in [79]. L. Console and P. Torasso have investigated several different versions of abductive diagnosis [22, 23], and have also developed an implementation of the theory as the CHECK system [120].

Even before the early rule-based, diagnostic systems mentioned above came into existence, diagnostic systems were built using probability theory [26, 37]. In particular, the simplified version of Bayes' rule, where observable findings are assumed to be conditionally independent given each of the possible disorders distinguished, and where disorders are assumed to be exhaustive and mutually exclusive, was and still is a popular basis for diagnostic systems [114]. The modelling power of this so-called 'independence form' of Bayes' rule is very limited, but it is known to offer a powerful basis for building classifiers, which can be adopted as a basis for diagnostic systems [28, 34]. However, the introduction of the Bayesian network formalism, also called Bayesian belief networks, causal probabilistic networks, or simply probabilistic networks [72, 52, 95], has rendered probability theory a viable alternative to logic and set theory for developing model-based diagnostic systems. A *Bayesian (belief) network*, also called causal probabilistic network, offer a model-based approach to the development of diagnostic systems, integrating qualitative knowledge and statistics. A Bayesian network consists of a directed acyclic graph (DAG) with a probability distribution over the collection of discrete random variables corresponding to the nodes in the DAG. The probability distribution may be factored according to the topology of the DAG, since it appears to be sufficient to specify only local probability distributions $\Pr(V|\pi(V))$, for each variable V with associated parents $\pi(V)$. This usually yields an enormous reduction in the amount of probabilistic information to be specified. The graph representation of a Bayesian network mirrors the dependencies and independencies among the variables corresponding to the nodes in the graph. A variable is dependent of its parents and children, but the variable V associated with a node is *conditionally independent* of any of its nondescendant nodes *given* the variables associated with its parent nodes. In this way the probabilistic information and also probabilistic computations are 'localised'.

As an example, consider the Bayesian network depicted in Figure 3, which models the genetics underlying Wilson's disease, also called hepatolenticular degeneration. This is a disease caused by a hereditary defect in the production of the copper-binding protein caeruloplasmin, which gives rise to precipitation of copper in many tissues, such as the liver, brain and the cornea (mainly in Descemet's membrane). The network provides a precise representation of the conditional independence assumptions in this domain. In Figure 3, for example, 'Free Serum Copper' is conditionally independent of 'Caeruloplasmin Serum Copper', given the value of 'Hepatic Copper'. The possibility to represent knowledge in a local, modular fashion, using the concept of causality as a structuring principle, is a very powerful feature of Bayesian networks. Like the model-based theories of diagnosis mentioned above, domain representation and diagnostic interpretation are separated in Bayesian-network theory

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 considera-




Figure 2: Basic idea underlying abductive diagnostic systems.

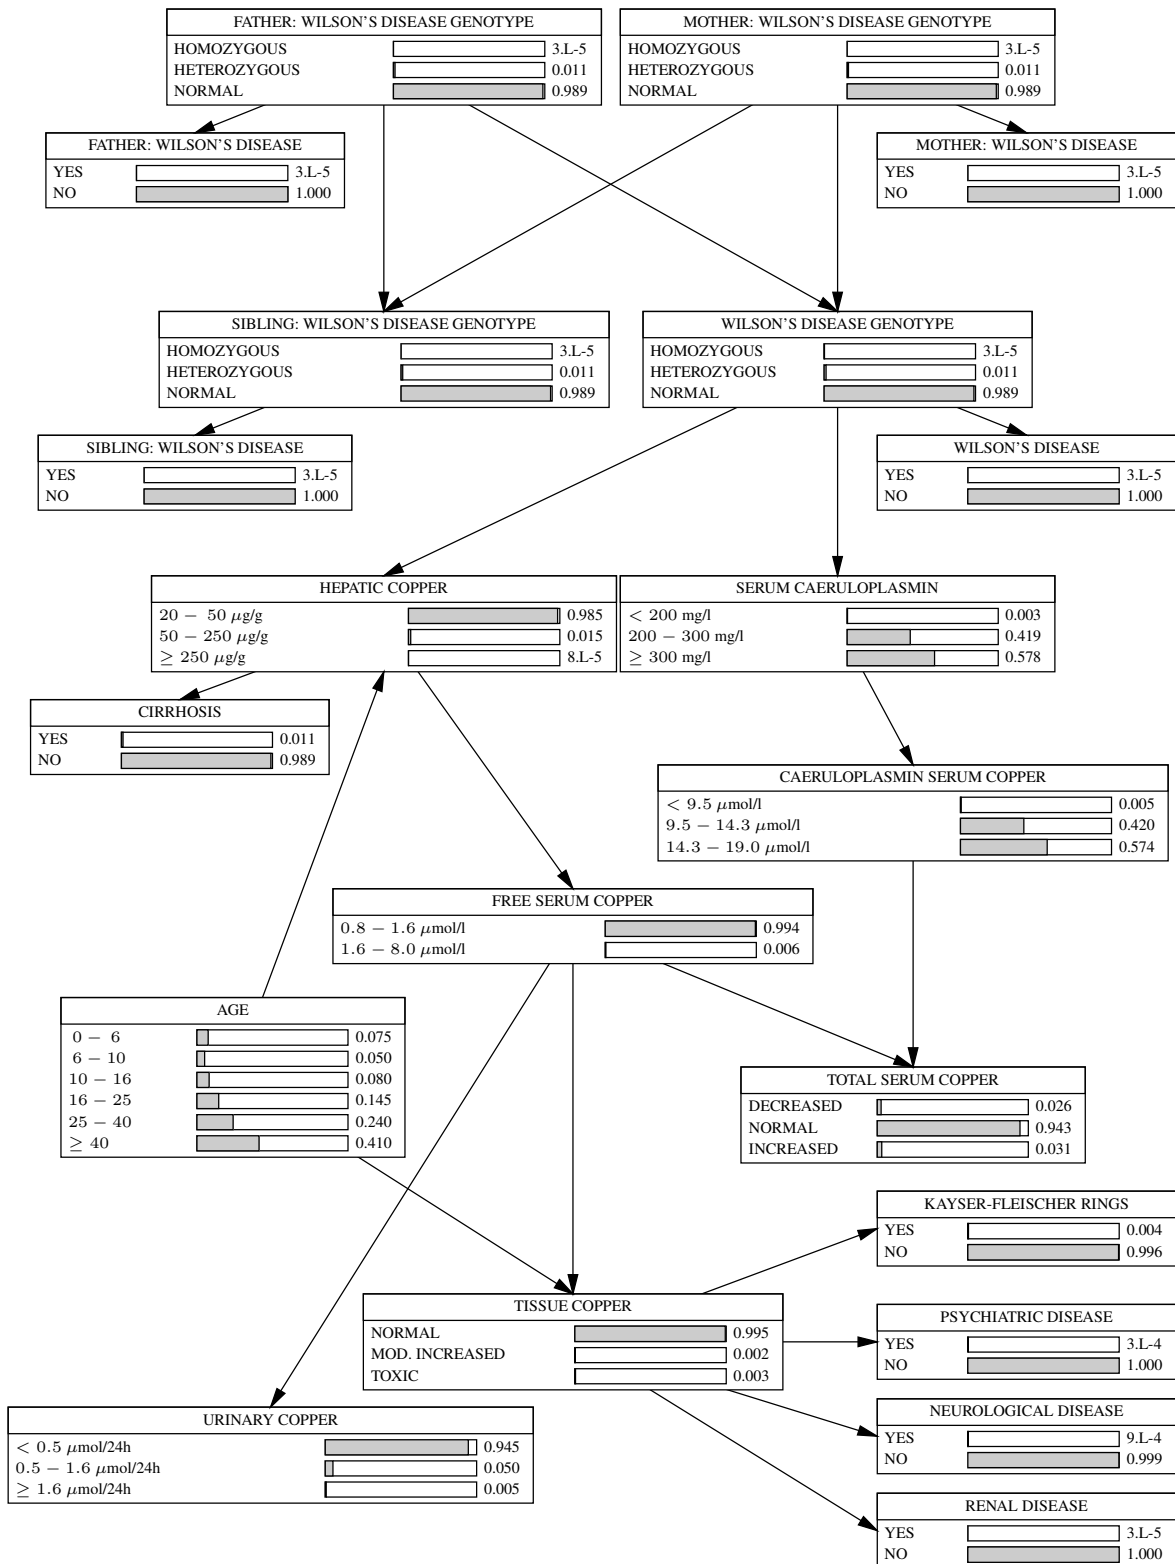


Figure 3: Bayesian network for the diagnosis of Wilson's disease [68].

tion 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 [58, 59] and Munin [7] systems.

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*, or MPA, to a given subset of variables. Establishing a maximum a posteriori assignment 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*. 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. The combination of the most likely values for these separate disease variables, however, need not be a maximum a posteriori assignment to these variables.

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 [8]. A test-selection 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.

The model-based approaches to diagnosis have a number of properties in common:

- It is possible to generate various collections of disorders as diagnoses, called *multiple-disorder diagnoses*, where the included disorders in a diagnosis together account for the observed findings.
- Methods for the sequential gathering of evidence during diagnostic reasoning, as have been traditionally available in rule-based diagnostic systems, can be added.
- A better understanding of generated advice by the user is obtainable, because the advice can be explained in terms of a model of the domain.

Finally, a number of formal frameworks of diagnosis have been proposed in recent years [24, 79, 117]. A framework of diagnosis offers sufficient generality to express and analyse various notions of diagnosis. These frameworks stress the important insight that there are many different notions of diagnosis possible, dependent on the nature of the domain. Some of these frameworks have been extended further in order to deal with temporal relationships in a domain [17]. For a detailed overview of techniques underlying symbolic model-based diagnosis, the reader is referred to [80].

2.1.3 The nature of medical domain models

The theories of diagnosis briefly discussed above, have no inherent relationships to medicine; these are general, domain-independent theories of diagnosis. Hence, only part of the story has been told; the other part concerns how medical knowledge fits in these theories.

In medicine several different types of knowledge play a part in clinical problem solving [78, 75, 98]. Due to the rich variety of types of medical knowledge, medicine has always been a suitable testbed for diagnostic reasoning research. In this section, we briefly describe the types of knowledge distinguished in the domain.

Medical textbooks like *Harrison's Principles of Internal Medicine*, [51], devote much space to describing the pathophysiological mechanisms underlying disease processes, i.e. disease processes are explained in terms of abnormal functional behaviour of organ systems. Sometimes, concise approximate mathematical models of normal or pathologically changed organ systems are available in the biomedical literature, but physicians generally prefer to discuss disorders in qualitative, often causal terms. As discussed above, the potentials of causal knowledge to act as a basis for diagnostic systems were already recognised in the context of the CASNET project [127]. The early probabilistic systems based on the independence form of Bayes' rule may also be viewed as encoding a very restricted type of causal knowledge. With the advent of abduction for qualitative causal models and Bayesian-network methods, neat formal methods for building diagnostic systems are now available for medical domains in which causal knowledge predominates.

Taxonomic principles also plays a role in the organisation of medical knowledge. Disorders are usually grouped as disorder categories; such hierarchical disorder organisations seem to help the clinician in coping with the enormous complexity of the medical domain. For example, in diagnosing jaundice of hepatobiliary origin, a clinician first tries to determine whether the disease affects the liver cells or the biliary tract. Assuming that the liver cells are damaged, the next step is to establish whether the disorder of the liver is acute or chronic in nature. It is well-known that such hierarchical structures may be quite effective for focusing problem solving in diagnostic systems [89].

Knowledge of the structure of the human body, i.e. anatomical knowledge, is of significant importance in a number of medical fields, such as surgery and neurology. Neurologists, for instance, use their in-depth knowledge of the structure of the peripheral and central nervous system in the diagnostic process all the time. LOCALIZE is an early diagnostic system that assisted in the diagnosis of damage to the peripheral nervous system [32]. In domains like neurology and surgery, knowledge of anatomical relationships will be almost always helpful in diagnostic problem solving [123]

Finally, time is also an important parameter in diagnostic problem solving in medicine. There are many disorders that are almost exclusively characterised by the time course of particular associated signs. A classical example is untreated lobar pneumonia, in which fever and other signs reach a climax, called the 'crisis', after several days, after which the condition of the patient suddenly improves. Nowadays, antibiotic treatment is instilled early in the disease process, and the typical temporal pattern is not longer observed. The notion of 'crisis' has even disappeared from recent medical textbooks. The time-dependent variations in levels of substances in blood are important sources of clinically relevant temporal patterns of disease, used in the modern process of diagnosis.

In model-based diagnosis in medicine, the types of knowledge mentioned above act as a basis for problem solving. Although empirical knowledge and clinical intuition remain extremely important in daily clinical practice, there is a general tendency in the medical community to delve deeper into the origins of disease. Hence, the increased interest of artificial intelligence researchers for model-based approaches is paralleled by similar interests in medicine as a whole.

2.2 Prognostic systems in medicine

2.2.1 Prognosis and outcome

In arguing about the trade-offs between costs and benefits of a particular medical procedure, knowledge of what is going to happen in the future with a patient and the associated environment is of major importance. Hence, patient management increasingly involves prediction of future events, i.e. *prognosis*. At the same time, as a consequence of the impact of Evidence-Based Medicine on daily practice, the physicians' tasks are gradually becoming more difficult. They are increasingly expected to be familiar with recent developments in their field of expertise, and to be able to balance the efficacy of novel tests and therapies against already established interventions. Obviously, here might lie a role for computer-aided decision support in the clinic. In fact, prognostic models already constitute, albeit sometimes implicitly, an integral part of a number of modern computer systems for diagnosis and treatment selection. In addition to their potential to improve the clinical management of the individual patient, prognostic models are used in a number of other ways as well. For example, a prognostic model may comprise a tool to assess the health state of patient groups, or be used to select patient cases for entering a clinical trial [91].

One of the main problems from the viewpoint of the model-based and qualitative reasoning practitioner concerns the suitability of a particular technique for building a prognostic model. Some thoughts are devoted to this important topic in the following.

2.2.2 Representation and construction of prognostic models

Various methods have been suggested for the representation of prognostic models, ranging from quantitative and probabilistic approaches to symbolic and qualitative ones. Although the notion of time seems to be a central issue in prognostic models, often time is only implicitly represented in order to keep such models manageable, both computationally and from the viewpoint of model construction.

Prognostic models can either be constructed by hand, with the help of human experts, or automatically, based on information extracted from databases. Often, the logical structure of a prognostic model is elucidated with the help of a domain expert, whereas underlying numerical information, like probabilistic information, is obtained from a database, yielding a semi-automatic model-construction method. Prognostic methods have been developed and studied both in the fields of decision analysis and statistics and in artificial intelligence (where we find the model-based and qualitative-reasoning approaches).

2.2.3 Traditional approaches: decision analysis and statistics

In the field of decision analysis it is assumed that *rational* decision making is accomplished according to the rules of mathematical decision theory, as developed in the 1940s by J. von Neumann and O. Morgenstern [125]. Decision theory is based on probability theory with utility theory; a rational decision or action is a decision that maximises expected utility, i.e. a utility in which uncertainties with respect to the outcomes of the decisions have been incorporated. Decision analysis is traditionally the field concerned with prognosis, because knowledge of possible future events is obviously important in selecting the best actions.

Decision trees The most popular tool to express sequential decision making is the *decision tree* (see the first four parts of the tutorial in [124]). Such a tree, with its root drawn to the left and leaves drawn to the right, offers an explicit representation of the relevant decisions and outcomes ordered according to progress of time. Every path in a decision tree from its root to a leaf, called

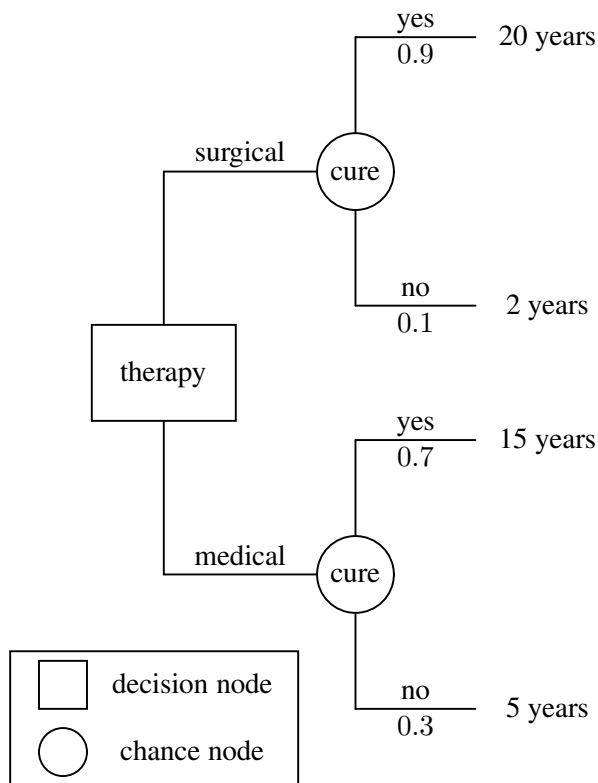


Figure 4: A simple decision tree.

a *strategy*, represents a potential course of events and actions. Because irrelevant strategies may be pruned according to the rules of decision theory, only strategies considered worth decision-analytic exploration need be represented (if irrelevant strategies were specified, they would be pruned anyway), which is a major advantage of the formalism.

Figure 4 provides an example of a decision tree that describes two possible actions (surgical or medical therapy) associated with the decision variable ‘therapy’ (rectangular node), with as outcome the chance variable ‘cure’ (circular node), denoting whether the patient is cured from the disease or not. Associated with a chance node is a probability distribution, conditioned on the variables on the path from the root to the associated node; for example, the number 0.9 is associated with the conditional probability $\Pr(\text{cure} = \text{yes} | \text{therapy} = \text{surgical})$. The leaves of the decision tree are labelled with utility function values, which in this case represent life expectancy. Following the rules of decision theory, it is possible to determine a therapy that maximises life expectancy (here the decision is to operate yielding life expectancy equal to 18.2 year) [113, 102].

Restricted probabilistic models A decision tree usually incorporates probabilistic information, but these probabilities would be typically derived from an external probability model, such as represented by Bayes’ rule or logistic regression equations. For example Figure 4 incorporates the probability $\Pr(\text{cure} = \text{yes} | \text{therapy} = \text{surgical})$, but if we would like to know the probability $\Pr(\text{therapy} = \text{surgical} | \text{cure} = \text{yes})$ (which percentage of patients who have been cured, have had surgery), the probabilities $\Pr(\text{cure} = \text{yes})$ and $\Pr(\text{therapy} = \text{surgical})$ would be required, because,

according to Bayes' rule,

$$\Pr(\textit{therapy} = \textit{surgical} \mid \textit{cure} = \textit{yes}) = \frac{\Pr(\textit{cure} = \textit{yes} \mid \textit{therapy} = \textit{surgical}) \Pr(\textit{therapy} = \textit{surgical})}{\Pr(\textit{cure} = \textit{yes})}$$

These probabilities would be available in an external probability model. The advantage of this is that when the structure of a decision tree changes, it is not always necessary to reconsult the database to compute new probabilistic information matching this new structure. Both Bayes' rule and logistic regression equations allow the interpretation of new information in light of current knowledge. In the decision-theoretic approach, such probabilistic models often reflect the prognostic part of a problem, but part may also be diagnostic in nature. A purely probabilistic approach may also be feasible in certain circumstances; then, logistic regression equations are still the models most popular for clinical decision support [38, 115]. For example, the APACHE III [66] scoring rule is used to predict life expectancy of intensive care patients based on the categories physiological measurements, age and previous health status. A score is given to each of these categories (based on a direct look-up table, for example an age above 75 corresponds to 6 points) and used in a linear equation to compute a summarising score. Logistic regression has proved its suitability in many domains but, it falls short once its assumptions, such as those concerning stochastic (in)dependencies or absence of induced dependencies among variables, are not met in practice. Logistic regression equations also rarely survive transfer to centres different from those where the data used for their construction were collected [114].

Markov processes Often it is difficult to calculate an expected outcome, for example calculating life expectancy when disease-specific mortality depends on age or when a decision may affect the complete future course of a patient's life. In these situations it is important to use an explicit representation of the history of a disease which includes the notion of time. A Markov process provides such a notion. Its representation consists of nodes, denoting states of the patient, and edges, denoting transitions between states within a period of time, called the *model cycle*, along with the transition probabilities. It is also possible to represent the effects of various actions on future states, together with immediate rewards when reaching a particular state, yielding a formalism that is known as Markov Decision Processes (MDP) [99]. When uncertainty is allowed with respect to the actual state of a system, for example because only part of the state is observable, an MDP is called a Partially Observable Markov Decision Process (POMDP) [74]. While in an MDP the problem is to find optimal sequences of actions given state transitions, in a POMDP optimal action sequences are sought with respect to probability distributions over the states.

Although the MDP formalism offers a very general and unrestrictive formalism for representing prognostic information including the effects of actions, it is assumed that the *Markov property* holds, i.e. it is assumed that the present state (of the patient or disease) no matter how it is arrived at, entirely determines the next state. Hence, only a small portion of the entire history is taken into account in decision making. This is not true for POMDPs, for which the entire history must be taken into account, which renders the formalism computationally expensive. An application of POMDPs to medical decision making is described in [42].

There are various algorithm in use in connection with Markov processes, varying from exact algebraic algorithms and dynamic programming algorithms to stochastic simulation algorithms [99].

As an example of a Markov process consider Figure 5. A patient in health state S_1 has probability of 0.95 to stay in S_1 after one year (assuming this is the model cycle), a probability of 0.04 to enter S_2 , and probability of 0.01 to die. This model can be easily simulated by a computer to calculate the life expectancy of say 100 patients and then calculating the average life expectancy. Moreover, one can modify the probabilities within a certain simulation to account for effects of ageing on mortality.




Figure 5: A simple Markov process model.

For example, if during a simulation a patient stays for 5 time units in S_2 (i.e. becomes five years older) one may increase the probability of death.

2.2.4 Model-based methods

In recent years several methods and techniques from the field of Artificial Intelligence have been introduced into patient management; in some of them prognostic knowledge is represented explicitly, whereas in others only implicitly. Usually, diagnostic systems do not contain any prognostic knowledge, although this simplification may not always be justified. Characteristics of prognostic models from the field of Artificial Intelligence include the *explicit representation* of domain models, which are often (partially) constructed by *automated learning*, and *automated reasoning* with the models.

Explicit representation of knowledge facilitates flexible reasoning and, when organised properly, supports knowledge reuse [1]. When domain knowledge is represented explicitly, emphasis is put on semantic concepts, such as notions of space, time and causality. Typical examples are models of human anatomy, (patho)physiology and biochemistry. Both anatomical and causal models have mainly been used in the field of medical diagnosis [75]. Anatomical knowledge can be very useful to localise causes of malfunction, in particular in neurology [20]. But such knowledge seems to have less relevance in the construction of prognostic models than causal knowledge, which additionally may also incorporate the notion of time explicitly. A typical example of the application of a qualitative prediction model for the purpose of diagnosis has been realised in the KARDIO system [15]. Basically, KARDIO's knowledge base consists of a logical formalisation of a qualitative simulation model of the (normal and abnormal) electrical activity of the heart. The predictive model can be triggered by the assumption of the presence of a particular (combination of) cardiac arrhythmias in the patient. This sets up a chain of events, finally leading to a collection of findings which represents an electrocardiogram (ECG), corresponding to the cardiac arrhythmias assumed to be present.

Where in a KARDIO-like approach, causally directed knowledge is used to obtain a predictive model, (patho)physiological models usually consist of algebraic equations. In such equations, some of the variables can be measured in the patient, and others must be derived from the known variables based on the equations. Such models can be applied in anaesthesiology in monitoring and controlling the physiological state of patients undergoing surgery. Qualitative reasoning is the AI field tradi-




Figure 6: General structure of a prognostic Bayesian network; each box denotes a part of the network.

tionally concerned with the specification of, and reasoning with, such models. For example, in the QSIM approach proposed by B.J. Kuipers, the time-dependent variation in the value of a variable is approximated by the notion of *qualitative state* [71]. Probabilistic AI models, for example based on Bayesian belief networks, also yield explicit representations of a domain. It often is possible to attach some causal interpretation to such models. They are very suitable for integration with traditional decision-theoretic methods (see below).

Model-based approaches as mentioned above have the potential to facilitate the development of knowledge-based systems, because the medical domain models can be (partially) obtained from the medical literature. When model construction is hard, less explicit domain models have been studied such as the use of case-based representations and its combination with more explicit domain models.

Unlike traditional statistical methods where the (mathematical) structure of models is given and probabilistic information is learnt from data, a characteristic of model-based approaches is that the structure of models can be learnt automatically. Machine learning (ML) is the field concerned with both learning structure and contents of (prognostic) models. Nevertheless, construction of the structure of prognostic models by hand is still the major development methodology. In ML, prognostic problems are often cast as classification problems. Automated learning usually implies a great deal of *search* in the hypothesis space defined by all possible models that attempt to explain the data given. Hypothesis spaces can be exponential in size and often infinite, therefore efficient ML techniques exploit an organisation of this space, often a hierarchical one, to focus search and avoid useless considerations of unattractive hypotheses.

A Bayesian network offers a model-based approach to the development of prognostic models, integrating qualitative, model-based approaches and statistics (See Figure 7). Bayesian networks are increasingly used in clinical medicine. As knowledge of the future is inherently uncertain, in prognostic reasoning uncertainty is even more predominant than in diagnostic reasoning. Another prominent feature of prognostic reasoning when compared to diagnostic reasoning is the exploitation of knowledge about the evolution of processes over time. Even if temporal knowledge is not represented explicitly, prognostic Bayesian networks still have a clear general temporal structure, which is depicted schematically in Figure 6. The outcome predicted for a specific patient is generally influenced by the particular sequence of treatment actions to be performed, which in turn may depend on the information that is available about the patient before the treatment is started. The outcome is often also influenced by progress of the underlying disease itself.

Formally, a prognosis may be defined as a probability distribution

$$\Pr(\text{Outcome} \mid \mathcal{E}, \mathcal{T})$$

where \mathcal{E} again is the available patient data, including symptoms, signs and test results, and \mathcal{T} denotes a selected sequence of treatment actions. The outcome of interest may be expressed by a single variable, for example modelling life expectancy. The outcome of interest, however, may be more

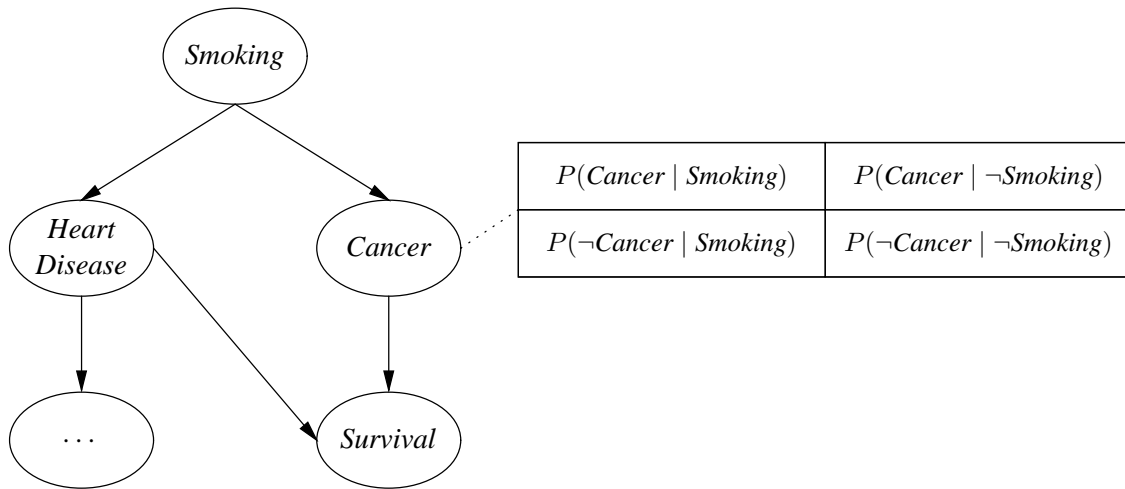


Figure 7: A prognostic Bayesian network that predicts the consequences of smoking.

complex, modelling not just length of life but also various aspects pertaining to quality of life. A subset of variables may then be used to express the outcome. An example is shown in Figure 7.

Prognostic Bayesian networks are a rather new development in medicine. Only recently have researchers started to develop such networks, for example in the areas of oncology [36, 77] and infectious disease [9, 82]. There is little experience as yet with integrating ideas from, for example, traditional survival analysis into Bayesian networks. Given the importance of prognostication in health-care, it is to be expected, however, that more prognostic networks will be developed in the near future.

2.2.5 Clinical usefulness

In the previous section, a large number of different techniques, all used in the development of prognostic models, have been discussed. Each of these techniques has its own virtues and limitations, and, obviously, it is not an easy matter to decide which technique is most suitable for developing a particular application. However, independent of the particular technique employed, there are a number of features a model should possess in order to qualify as being potentially clinically useful, although these features have no universal significance. A number of these features will be discussed below.

The development of an accurate prognostic model usually takes a great deal of time. With respect to the universality of models, it would, therefore, be a waste of time and effort if such models could not be transferred to other clinical centres than the centre for which the model was developed. Prognostic models that may survive such transfer are called *robust*: slight changes to the environment in which a prognostic model is employed, or slight changes to its contents, will not have a major effect on its level of performance. Techniques that allow incorporating explicit domain knowledge frequently yield models that are more robust, and thus can be transferred to other centres of clinical care. Prognostic models should also be easy to use and understand by clinicians. Again, model-based techniques tend to have this virtue, although constructing explicit domain models may be very time consuming, which is a major disadvantage of model-based techniques. Logistic regression equations do not give any insight into underlying disease mechanisms, but are still very popular in medical research, possibly due to their simplicity and the long period of time they have been around. Of course, variables that are routinely available, or easily obtainable, should be included in a prognostic model if relevant, and this may also promote the use of explicit models.

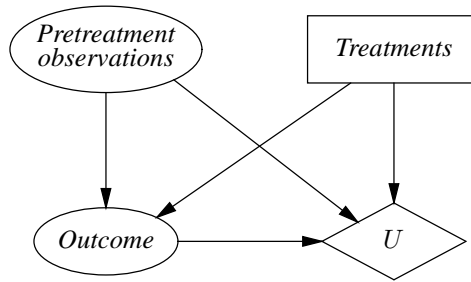


Figure 8: General structure of an influence diagram, including a prognostic Bayesian network and a utility node U ; each ellipse and box denotes a part of the diagram.

Prognostic modelling has a long-standing tradition in the fields of decision analysis and statistics, and an increasing number of specific prognostic models is being developed in Medicine. On the other hand, it has been noticed that many of these prognostic models are seldom used in clinical practice [130]. A possible explanation for this phenomenon may be that many prognostic models have little semantic relationship with the clinical domain. A model-based approach to prediction and prognosis, as advocated in AI, may offer advantages in this respect, by offering more natural and intuitively attractive tools for decision support.

The construction of prognostic models is usually carried out using information extracted from clinical databases. Therefore, research in particular in machine learning seems to provide opportunities for automatically learning these expressive models. Unfortunately, clinical databases of sufficient size and reliability, filled with data from studies with rigorous patient follow-up, are still not widely available.

2.3 Treatment selection in medicine

The formalism of Bayesian networks provides only for capturing a set of random variables and a joint probability distribution over them. A Bayesian network therefore allows only for probabilistic reasoning, as in establishing a diagnosis for a specific patient and in making a prediction of the effects of treatment. For making decisions, as in deciding upon the most appropriate treatment alternative for a specific patient, the network formalism does not provide. Reasoning about treatment alternatives, however, involves reasoning about the effects to be expected from the different alternatives. It thus involves diagnostic reasoning and, even more prominently, prognostic reasoning. To provide for selecting an optimal treatment, a Bayesian network and its associated reasoning algorithms are therefore often embedded in a decision-support system that offers the necessary constructs from decision theory to select an optimal treatment given the predictions [9, 77]. Alternatively, the Bayesian-network formalism can be extended to include knowledge about decisions and preferences. An example of such an extended formalism is the influence-diagram formalism [108]. Like a Bayesian network, an *influence diagram* includes an acyclic directed graph. In this graph, the set of nodes is partitioned into a set of probabilistic nodes modelling random variables, a set of decision nodes modelling the various different treatment alternatives, and a value node modelling the preferences involved. Influence diagrams for treatment selection once again have a clear general structure, which is depicted schematically in Figure 8.

2.4 Qualitative modelling and simulation in bioinformatics

2.4.1 Challenges to molecular biology

One of the remarkable characteristics of molecular biology today is the upscaling of its experimental methods to the genomic level. Hardly imaginable only 20 years ago, the sequencing of complete genomes has become a routine job, highly automated and realised in a quasi-industrial environment [41]. The miniaturisation of techniques for the hybridisation of labelled nucleic acids in solution to DNA molecules attached to a surface has given rise to DNA chips, tools for measuring the level of gene expression in a massively parallel way [73]. The development of proteomic methods based on two-dimensional electrophoresis, mass spectrometry, and the double-hybrid system allows the identification of proteins and their interactions at the genomic level [92].

These new experimental methods generate enormous volumes of data on different aspects of the cell. On the one hand, they identify the interactions between genes, proteins, metabolites, and other molecules, defining the structure of biochemical networks. On the other hand, they measure the evolution of the state of the cell, that is, the temporal variations in the concentration and the localisation of the molecular components as a function of changes in environmental conditions. The major challenge for what is called *functional* or *systems biology* consists in relating these structural and behavioural data in order to obtain a global interpretation of the functioning of an organism [64]. The aim is to predict and understand how the observed behaviour of an organism – the adaptation to its environment, the differentiation of cells during development, or even the evolution of the species in the long run – can emerge from the whole of molecular interactions.

In addition to high-throughput experimental methods, mathematical and bioinformatical approaches are indispensable for the analysis of biochemical networks. Given the large number of components of most networks of biological interest, connected by positive and negative feedback loops inside and between cells, an intuitive comprehension of the spatiotemporal evolution of a system is often difficult, if not impossible to obtain. Mathematical modelling supported by computer tools can contribute to the analysis of a biochemical network by allowing the biologist to focus on a restricted number of plausible hypotheses. The formulation of a mathematical model requires an explicit and non-ambiguous description of the hypotheses being made on the regulatory mechanisms under study. Furthermore, through analysis or simulation, the model yields predictions on the behaviour of the cell that can be verified experimentally.

A variety of methods for the *modelling and simulation of genetic, metabolic, and signal transduction networks* have been proposed in the literature [53, 30, 39, 44, 84, 111]. The use of formal methods to study biochemical networks is currently subject to two major constraints. First of all, the reaction mechanisms underlying the interactions are usually not or incompletely known. This prevents the formulation of detailed kinetic models, such as those developed for the genetic switch controlling phage λ growth [85] or the feedback mechanisms regulating tryptophan synthesis in *E. coli* [106]. A second constraint arises from the general absence of quantitative information on most kinetic parameters and molecular concentrations. As a consequence, traditional methods for numerical analysis are difficult to apply.

2.4.2 Qualitative reasoning and simulation

The above constraints have stimulated an interest in the modelling and simulation methods developed within *qualitative reasoning (QR)* [128, 31], the first example probably being the qualitative models of gene expression developed by Koile and Overton [67]. The basic idea of qualitative simulation consists in the abstraction of a discrete description of a biochemical network from a continuous model,

followed by the analysis of the discrete instead of continuous equations to draw conclusions about the dynamics of the system. One of the best-known formalisms developed for this purpose are the *qualitative differential equations (QDEs)* used in the simulation method QSIM [69, 71]. QDEs are abstractions of ODEs of the form

$$\frac{dx_i}{dt} = f_i(\mathbf{x}), 1 \leq i \leq n, \quad (1)$$

where $\mathbf{x} = [x_1, \dots, x_n]^T$ and $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is a linear or nonlinear function. The variables x take a *qualitative value* composed of a qualitative magnitude and direction. The qualitative magnitude of a variable x_i is a discrete abstraction of its real value, while the qualitative direction is the sign of its derivative. The function f_i is abstracted into a set of *qualitative constraints* which restrict the possible qualitative values of the variables. Given an initial *qualitative state*, consisting of the qualitative values for \mathbf{x} at the initial time-point, the QSIM algorithm generates a tree of *qualitative behaviours*. Each behaviour in the tree describes a possible sequence of state transitions from the initial state. It has been proven that every qualitatively distinct behaviour of the ODE corresponds to a behaviour in the tree generated from the QDE, although the reverse may not be true.

The incomplete understanding of many genetic regulatory mechanisms on the molecular level, and the general absence of quantitative knowledge, has stimulated an interest in qualitative simulation techniques. An example of their application is given by Heidtke and Schulze-Kremer [43]. Their case-study concerns λ phage growth control in *E. coli*. The model consists of seven QDEs, representing the different stages of λ phage growth as it follows either the lytic or lysogenic pathway. Each QDE describes in detail the infected bacterium, including explicit representation of λ phages inside and outside the cell, viral DNA, ribosomes, mRNA, and proteins, and the way in which they interact to control major cellular events. For other examples of the application of qualitative reasoning methods to the analysis of biochemical networks, see [6, 47, 122, 121]

2.4.3 Bayesian networks and the discovery of functional interactions

Bayesian networks are also being explored to bioinformatics applications. The insight obtained by constructing Bayesian networks, in particular when done automatically by using one of the learning methods, may yield scientifically useful results. As the topology of a Bayesian network can be interpreted as a representation of the uncertain interactions among variables, there is a growing interest in bioinformatics to use Bayesian network for the unravelling of molecular mechanisms at the cellular level. For example, finding interactions between genes based on experimentally obtained expression data in microarrays is currently a significant research topic [35]. Biological data are often collected over time; the analysis of the temporal patterns may reveal how the variables interact as a function of time. This is a typical task undertaken in molecular biology. Bayesian networks are now also being used for the analysis of such biological time series data [103].

2.4.4 Upscalability

A problem with existing qualitative simulation approaches is their limited upscalability. As a consequence of the weak nature of qualitative constraints, and the difficulty to identify implicit constraints, behaviour trees quickly grow out of bounds. This causes the range of application of the methods to be limited to regulatory systems of modest size and complexity. Systems of even a few elements related by positive and negative feedback loops cannot be handled, unless these systems have been



Figure 9: Series C pacemaker.

so well-studied already that behaviour predictions can be tightly constrained. Since the early 1970s, biomathematicians and theoretical biologists have studied a class of piecewise-linear linear differential equation models of genetic networks that put strong constraints on the (local) dynamics of the system [57, 86, 112, 119]. By reframing the mathematical analysis of these models in a qualitative simulation framework, using a simulation algorithm tailored to the equations, larger networks with complex feedback loops can be treated [55, 56]. The resulting method has been implemented in the computer tool Genetic Network Analyzer (GNA), and has been applied to the simulation of the initiation of sporulation in *B. subtilis* [54, 40].

Extensions of qualitative simulation methods allow weak numerical information in the form of interval bounds on the qualitative magnitude and direction of the variables, and numerical envelopes around the functions f_i , to be integrated in the simulation process [14]. This presents another way to restrict the number of behaviours, while it may also refine the predictions of the remaining behaviours. Heidtke and Schulze-Kremer [43] show how the conclusions of their model can be made more precise by adding semi-quantitative information.

Much work in qualitative reasoning has focused on the automated composition of an appropriate model of a system, given a user query, from a situation description and a domain knowledge base [63]. One of the basic approaches towards automated modelling, *qualitative process theory* [33], has been used by Karp in his method for the construction and revision of gene regulation models [60, 61, 62]. From a taxonomic knowledge base describing biological objects like strains of bacteria, genes, enzymes, and amino acids, and a process knowledge base comprising a theory about biochemical reactions, the program GENSIM predicts the outcome of a proposed experiment. If the predictions do not match with the observations made while actually carrying out the experiment, then the program HYPGENE generates hypotheses to explain the discrepancies. In particular, it revises assumptions about the experimental conditions that GENSIM used to derive its predictions. The programs have been able to partially reproduce the experimental reasoning underlying the discovery of the attenuation mechanism regulating the synthesis of tryptophan in *E. coli*.

3 Model-based and qualitative reasoning methods in biomedical industry and health-care

In this section, we describe a number of model-based and qualitative reasoning solutions that have been developed in the context of biomedical research and health-care. It is not the purpose of this



Figure 10: Pacemaker programmer.

section to be exhaustive, but rather to convey an impression about to what extent model-based and qualitative-reasoning techniques are being explored in these areas. Selection of examples has mainly been based on the availability of detailed information concerning the examples.

3.1 Reprogramming pacemakers

Vitatron, an international firm devoted to the production of heart pacemakers with headquarters located in the Netherlands, has been working on model-based computer systems intended to assist cardiologists in programming heart pacemakers.

A pacemaker consists of a *can*, which contains a microprocessor, a battery and an impulse generator. Impulses are transmitted to the heart by means of a *lead*, which is attached to the can's connector. A lead is either unipolar or bipolar; a *unipolar* lead contains one insulated coil, whereas a *bipolar* lead contains two coils, separated by an inner insulation. An outer insulation shields a lead from the environment. The tip of a lead, which contains an electrode, is implanted into the inner, endocardial surface of the heart; the actual location depends on the type of pacemaker. The pacemaker can is usually implanted in the pectoral region, with the lead running through the right subclavian vein to the internal surface of the heart. An ample of a modern, advanced pacemaker is Vitatron's *series C* pacemaker (See Figure 9).

A pacemaker is programmed by means of a *programmer*, a computer with a special user interface for data entry and display, and with special software to control an external magnet that communicates with the pacemaker. Figure 10 shows an example of a pacemaker programmer. The magnet is placed above the location of the pacemaker; information from the programmer to the pacemaker, and back, is transmitted by means of telemetry.

Researchers of Vitatron together with researchers now at the University of Nijmegen have constructed an extensive causal model that integrates pacemaker malfunctioning of the Diamond II pacemaker and resulting signs and symptoms in the patient [81]. The purpose of this system was to deliver decision support to cardiologists who wish to (re)programme a patient's pacemaker. Using abductive diagnosis, as described in Section 2.1.2, as a basic techniques a prototype system has been developed using David Poole's Theorist system [97]. A simpler version of the system was introduced in 2003 for use in conjunction with the recently introduced, more advanced in comparison to the Diamond II,

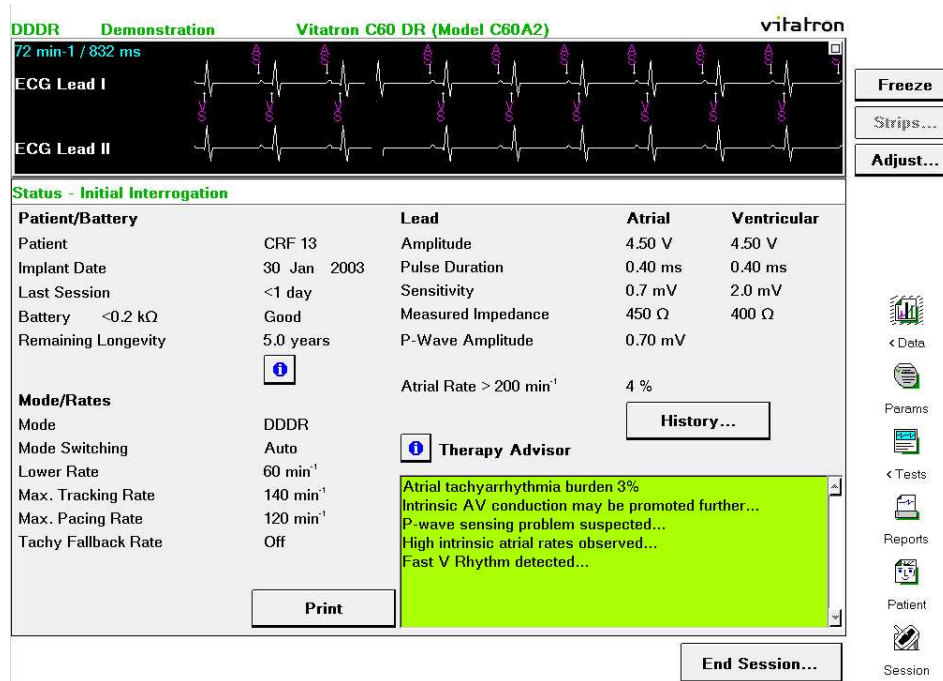


Figure 11: Therapy Advisor start-up screen with status information.

series C pacemaker [83]. This system is called the *Therapy Advisor*.

The Therapy Advisor is an integral part of the pacemaker-programmer (cf. Figure 10) software that is used at the pacemaker follow-up procedure. The first step of a follow-up that involves the programmer is the acquisition of pacemaker data. Secondly, the Therapy Advisor is activated after which the start-up screen, called the status screen, is presented to the user (cf. Figure 11). It displays the lead recordings of an ECG, the overall status of the pacemaker; at the lower part of the screen, messages generated by the Therapy Advisor are presented. The messages shown here are general in nature. A distinction is made between the following types of messages:

- pacemaker technical problems that result in the entire loss of diagnostics, e.g. pacemaker reset and battery depletion;
- pacemaker-patient interface problems. These problems arise when the technical pacemaker parameters, e.g. sensitivity or rate response, are not correctly aligned with the patient's needs and usually render parts of the diagnostics unreliable, e.g. far field R-waves distorted atrial sensing;
- clinical causes, possibly related to complaints or deterioration (e.g. atrial fibrillation and retrograde conduction);
- clinical effects, possibly related to complaints or deterioration (e.g. fast ventricular rhythm and low AV synchrony).

A separate screen is available to the user for more detail about the detected irregularities. This detailed information can consist of an explanation of the diagnostic observations, characterisation of the clinical cause (e.g. atrial fibrillation), therapy parameter advice, technical parameter advice and diagnostic parameter advice (cf. Figure 12).

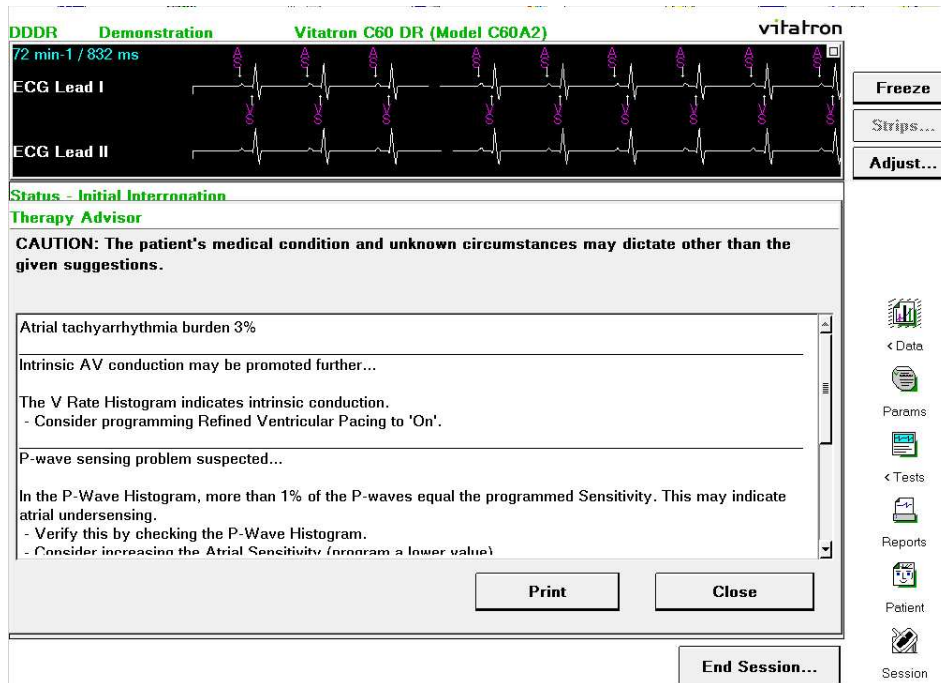


Figure 12: Therapy Advisor screen with advice.

The inference algorithm treats the model underlying the system as a forest of trees, where each tree can be interpreted as part of the definition of an evidence function e as discussed above. During problem solving, a search algorithm, implemented in C++, traverses the trees sequentially using the data from the pacemaker. As the main motivation for incorporating the Therapy Advisor into the pacemaker programmer software was to reduce time needed for follow-up, all data interpreted by the system come from the pacemaker. This also implies that information that is not available to the pacemaker, e.g. information obtained from the medical interview with the patient, is currently not taken into account by the Therapy Advisor in generating recommendations.

3.2 Challenging QR research issues in electrocardiology

Electrocardiology is a stimulating and promising application domain for qualitative and model-based reasoning research. Much of the research effort has been so far aimed at the automated interpretation of electrocardiograms (ECGs), for which the interpretative rationale is well established. However, a major drawback of ECGs, which record electrical signals from nine sites only on the body surface, is their poor spatial coverage. Some arrhythmias, in addition to conduction anomalies caused by infarcts and ischaemia, can be missed by an ECG.

Thanks to the latest advances in technology, a far more informative technique, namely body surface mapping (BSM), is becoming more widely available. In BSM, the electric potential is measured over the entire chest surface, and more spatial information about the heart electric activity is available, though the ability to relate visual features to the underlying complex phenomena still belongs to very few experts [116]. Besides body surface maps, also epicardial or endocardial maps are becoming available, either measured or obtained non invasively from body surface data through mathematical inverse procedures. These kinds of maps are the most precise in locating the anomalous conduction sites [21].

While a rationale for the interpretation of electrocardiographic maps is progressively being defined, the goal of bringing such techniques into the clinical practice is getting closer. In this context, an important role 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 help him in the definition of an interpretative rationale, and (2) achieve the long-term goal of completely automating map interpretation to be used in a clinical context. The delivery of such a tool would be of great impact on health care as the interpretation of cardiac potential maps gathered by endocavitary probes is essential for the localisation of the ectopic sites associated with ventricular arrhythmia.

At the current status of research, map interpretation mainly deals with activation isochronous maps and sequences of isopotential maps. The former ones deliver a lot of information about the wavefront propagation: each contour line aggregates all and none but the points that share the same excitation state, and subsequent increasing isocurves are nothing but subsequent snapshots of the travelling wavefront; whereas the latter ones give temporal information about electric phenomena, such as current inflows and outflows. The interpretation of activation isochronous maps currently grounds on features such as contour shapes, the directions along which the front velocity is higher, minimum and maximum regions. The contour shapes and highest velocity directions reflect the conduction anisotropy, and the fibre architecture. Minimum and maximum correspond to the points where the front emerges and extinguishes, respectively. Such features can be correlated to the expected activation sequence, and used in a diagnostic context: deviations from the expected patterns would highlight ectopic sites associated with ventricular arrhythmias.

Conventional contouring tools would allow us to efficiently visualise patterns of electric potential distribution but they do not facilitate the automated extraction of general rules to infer the correlation between pathophysiological patterns and wavefront structure and propagation. The Spatial Aggregation (SA) approach [10, 46], aiming at interpreting a numeric input field, is potentially capable of capturing structural information about the underlying physical phenomenon, to identify its global patterns and the causal relations between events. These characteristics are due to its hierarchical strategy in aggregating objects at higher and higher levels. However, when dealing with complex domain geometry and/or with non uniform data meshes, as in this context, the original version of SA may unsoundly perform contouring: isocurve entanglement and/or segmentation phenomena may occur due to metric-based adjacency relations and a scarce density of isopoints.

As robust contouring is a must in view of the identification and extraction of salient features and structural information about the underlying electrical phenomenon, a thorough revision and extension of SA is necessary to deal with the 3D geometries of heart and chest, where data are given on the nodes of non uniform quadrangular meshes: either planar 2D elements when transversal or longitudinal sections of the myocardium are considered, or 3D surface elements when the epicardial surface is explored.

Some work has still to be done to improve the robustness of SA as to the construction of isocurves from a non-uniform numeric input field [49, 50] by providing new algorithms and definitions for soundly building neighbourhood relations upon 2D complex domain geometry. This is a meaningful result from our application point of view as 3D processes are suitably studied also by transversal/longitudinal sections of the 3D domain. It will allow us to identify and define spatial relations between isocurves, as well as between possible higher-level aggregated objects. Such relations should reveal adjacency and interactions between the spatial objects, and, then, underlie the feature extraction process.

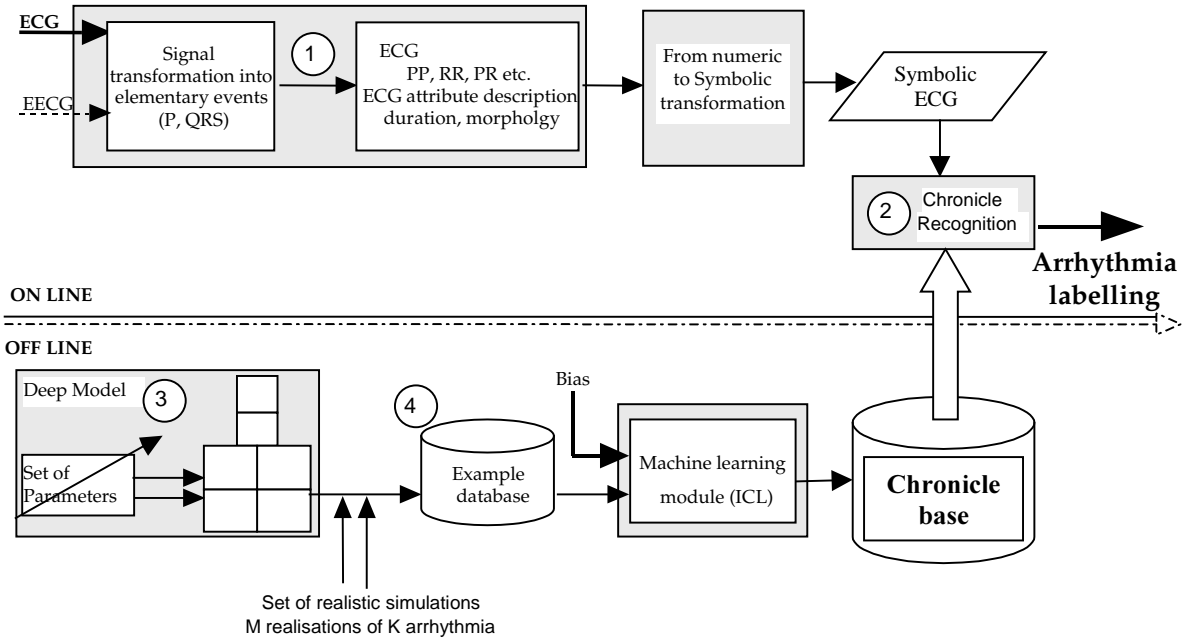


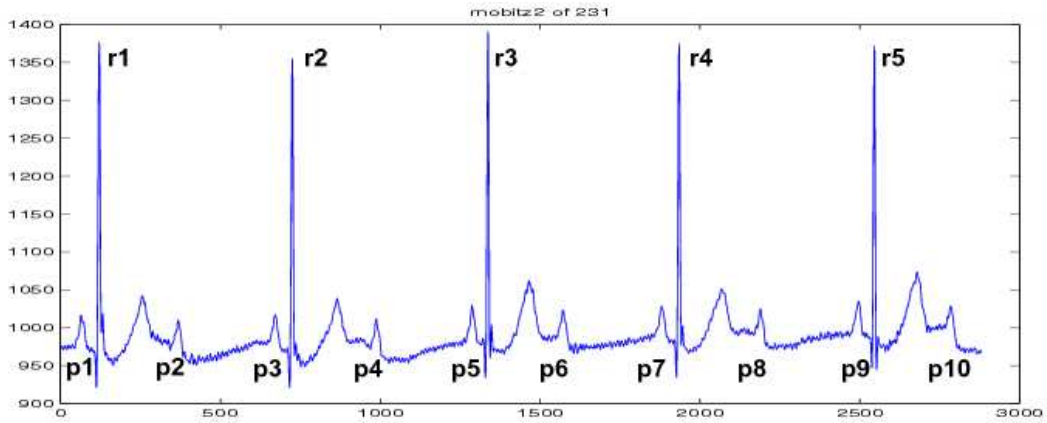
Figure 13: Architecture of the described approach.

3.3 A model-based approach for learning to identify cardiac arrhythmias

The aim of this work was to provide an integrated framework enabling the input signal analysis, to guarantee early detection of abnormal situations and to provide helpful characterisations of arrhythmia disorders by relying on medical knowledge [18]. The main contribution is to integrate temporal abstraction and diagnosis techniques. The approach can be described in two steps (see Figure 13): the first one is done off-line and the aim is to build a set of high-level symbolic characterisations of the cardiac arrhythmias [100, 126]. The second one is an on-line step which is in charge of analysing the signal and identifying arrhythmias by matching the symbolic representation of the signal to prestored characterisations.

The off-line step relies on automatic learning techniques and its aim is to extract the most discriminating patterns which can identify normal and pathological ECGs. These patterns, named chronicles, are sets of temporally constrained events. The learning step relies on inductive logic programming (ILP) techniques which have led to the discovery of interesting knowledge in several scientific domains such as biology or (pharmaceutical) chemistry [90]. The tool chosen for experimentation is the ICL system [101]. It makes use of learning examples (see number 4, Fig. 13) which are either real signals (like the labelled ECG signals from the MIT-BIH database [88]) or signals obtained by simulating arrhythmias on the Carmen [45] cardiac model (number 3, Figure 13).

The on-line part includes the temporal abstraction module (number 1, Figure 13), based mainly on signal processing and classification tools, which outputs the symbolic representation of the analysed signals in terms of time-stamped events, and the chronicle recognition module (number 2 in Figure 13) which analyses the stream of events and identifies arrhythmia disorders by detecting characteristic patterns. The temporal abstraction module is in charge of analysing multiple ECG signals. Chronicle recognition consists in skimming the flow of events coming from an observed process and detecting the specific events that belong to a chronicle [29]. This process is very similar to pattern-matching



```

begin(model(mobitz2_231_1)).
mobitz2.
  p_wave(p1, 124, normal, noPreWave).
  qrs_complex(r1, 295, normal, p1).
  p_wave(p2, 968, normal, r1).
  p_wave(p3, 1812, normal, p2).
  qrs_complex(r2, 1967, normal, p3).
  p_wave(p4, 2687, normal, r2).
  p_wave(p5, 3523, normal, p4).
  qrs_complex(r3, 3669, normal, p5).
  p_wave(p6, 4306, normal, r3).
  p_wave(p7, 5159, normal, p6).
  qrs_complex(r4, 5335, normal, p7).
  p_wave(p8, 6030, normal, r4).
  p_wave(p9, 6869, normal, p8).
  qrs_complex(r5, 7034, normal, p9).
end(model(mobitz2_231_1)).

```

Figure 14: A Mobitz type II arrhythmia ECG and its description as an ICL example.

```

class(bigeminy)      :- % [15, 0, 0, 0, 0]
  qrs(R0, abnormal, _),
  p_wav(P1, normal, R0), qrs(R1, normal, P1),
  qrs(R2, abnormal, R1), rr1(R1, R2, short).
class(bigeminy)      :- %[5, 0, 0, 0, 0]
  qrs(R0, normal, _),
  p_wav(P1, normal, R0), qrs(R1, abnormal, P1).
class(lbbb)          :- % [0, 20, 0, 0, 0]
  qrs(R0, abnormal, _),
  p_wav(P1, normal, R0), qrs(R1, abnormal, P1).
class(mobitz2)       :-% [0, 0, 17, 0, 0]
  p_wav(P0, normal, _), equal(P0, R0),
  p_wav(P1, normal, R0), qrs(R1, normal, P1).
class(mobitz2)       :-,%[0, 0, 3, 0, 0]
  p_wav(P0, normal, _), equal(P0, R0),
  p_wav(P1, normal, R0), qrs(R1, abnormal, P1).
class(pvc)           :-%[0, 0, 0, 0, 20]
  p_wav(P0, normal, _), qrs(R0, normal, P0),
  p_wav(P1, normal, R0), qrs(R1, normal, P1),
  qrs(R2, abnormal, R1), rr1(R1, R2, normal).
class(normal)        :-% [0, 0, 0, 19, 4]
  p_wav(P0, normal, _), qrs(R0, normal, P0),
  p_wav(P1, normal, R0), qrs(R1, normal, P1),
  p_wav(P2, normal, R1), qrs(R2, normal, P2),
  p_wav(P3, normal, R2), qrs(R3, normal, P3),
  p_wav(P4, normal, R3).

```

Figure 15: Examples of rules learned with ICL for five classes of arrhythmias.

associated with temporal constraint satisfaction.

Some new issues are currently investigated. Firstly, the work is extended to cope with multiple sources of information (multichannels and multisensors) in order to improve the robustness of signal detection and especially P wave detection. Secondly, it appears important to study a hierarchical recognition task. This consists in relying on the easily detectable subparts of the signal (like the QRS complexes) as long as they are sufficient and triggering a focussed analysis (to detect the P wave for instance) only when it is needed. It means developing a strong connection between the temporal abstraction module and the chronicle recognition tool, the second driving the first.

The system was initially developed in order to improve early arrhythmia detection in the ICCU context but concerns several application domains. For instance, it can be interesting in the context of external ambulatory systems (Holters) as well as that of implantable cardiac devices. In this latter case, the fact that the system can adapt to the characteristics of each patient by learning the detection thresholds, for instance, is clearly a very positive point. Another important feature is the automatic acquisition of expertise. Due to technological advances in signal acquisition (new sensors for instance), clinicians have to rapidly update their current expertise to efficiently interpret these new data. The approach does not rely on this human expertise and the learning module is well suited to acquire a new set of signatures on these new signals. For instance, it could be applied to active cardiac devices which have recently been proposed and which rely on leads located in both ventricles. These new devices can tackle both rhythmic and hemodynamic disorders and the signatures, which are still poorly known, can be acquired by an automatic learning method.

3.4 Qualitative reasoning for the modelling of metabolic systems

The prediction of the evolution over time of the patient's state plays a crucial role both in a diagnostic and therapeutic medical context. A traditional way to approach such a problem deals with both the formulation of mathematical models of the dynamics of patho-physiological systems and the simulation of their behaviour [19]. Such models, which are generally described by Ordinary Differential Equations (ODE), are computationally tractable with classical methods which allow us to derive, either analytically or numerically, meaningful predictions of the behaviour of the considered system. But, for the medical domain as for many other physical domains, quantitative model formulation may be not successfully applicable due to the incompleteness of the available knowledge about either the functional relationships between variables or the numerical values of model parameters, which could be non identifiable both for the lack of adequate experimental settings and for the impossibility of measuring 'in vivo' the values of a few variables. Qualitative modeling methods are capable to cope with difficulties in model building in presence of incomplete knowledge [71]. But, whereas qualitative predictions of a patho-physiological system behaviour may be properly exploited in the testing phase of diagnostic reasoning [48], they are almost always inadequate to be used in a therapy planning context as the effects of different therapies are requested to be deeply investigated at a quantitative level.

As alternative to conventional mathematical modeling frameworks, the so-called input-output approaches, that are able to describe the dynamics of a real system from input-output data, have been proposed. Neural networks, multi-variate splines and fuzzy logic systems are the best known approximation schemes used for learning an input-output relation from data. Although these approaches are successfully applied to a variety of domains, they are affected by two main drawbacks that are particularly serious in medicine: first, the identification result, a non-linear function, does not capture any structural knowledge; second, the model identification procedure usually requires a large amount of data and is often extremely inefficient. In practice, such methods fail when the experimental data set is poor either in size or in quality. Such a situation is not rare in the case of metabolic systems as they are very often characterized by an intrinsic difficulty in performing experiments and in measuring the variables of interest.

Qualitative Reasoning may effectively be integrated with classical input-output approaches to solve the problems above in a great deal of situations: pathophysiological knowledge is very often available even if insufficient to formulate a quantitative model, and it could be conveniently embedded into a fuzzy identifier. FS-QM [11] is a hybrid approach which is half way between the structural and input-output approach. It uses the outcomes of the simulation of a qualitative model to build a good initialization of a fuzzy system identifier. Such an initialization allows us to efficiently cope with both the incompleteness of knowledge and the inadequacy of the available data set, and to derive an accurate input-output model of a great deal of metabolic systems also in data poor contexts where conventional methods fail.

The resulting model which embeds pathophysiological knowledge provide an interpretative key of the underlying mechanisms. The range of applicability of FS-QM to study metabolic systems is quite large as shown by its application to study a dynamic pathological system in response to exogenous perturbations, namely the blood glucose level in insulin-dependent diabetes mellitus patients in response to insulin therapy and meal ingestion [13], and to successfully identify the nonlinear dynamics of intracellular thiamine in the intestine tissue [12]. Let us observe that the classical compartmental approach to metabolic system modeling revealed to be inapplicable to model the latter system for the incompleteness of the available knowledge and for the difficulty of gathering an experimental data set rich enough to guarantee the well-posedness of the parameter estimation procedure. The poor data set

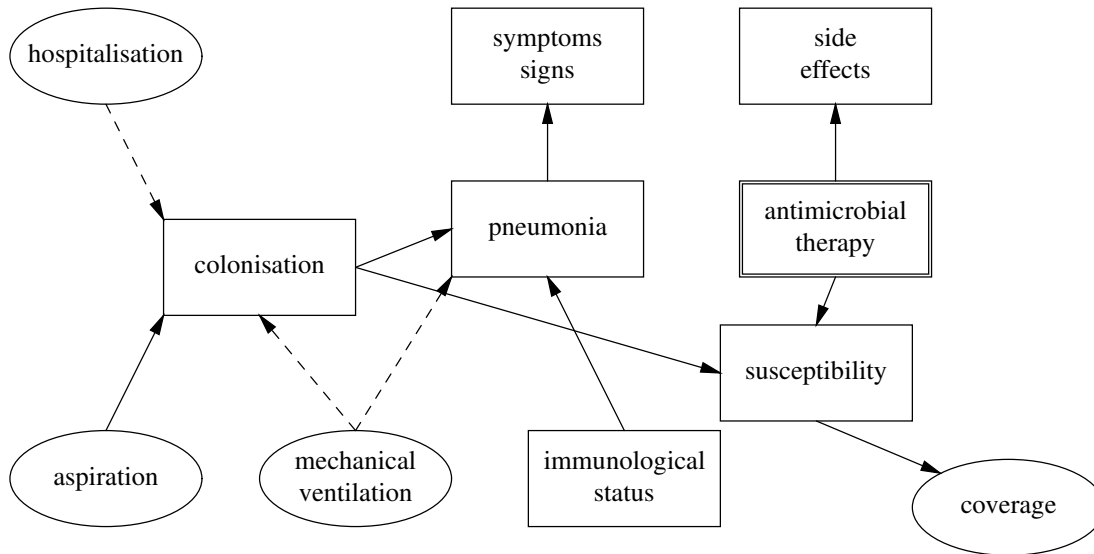


Figure 16: Global structure of the Bayesian-network model of VAP.

has been also responsible of the failure of input-output nonlinear regression approaches.

The quantitative assessment of different pathophysiological processes is often a necessary step for an insightful interpretation of metabolic systems and for therapy planning. The integration of QR techniques with conventional modeling approaches plays a key role towards the development of robust and efficient quantitative modeling methods capable to deal with the complexity of the biomedical systems.

3.5 Advice in the prescription of antibiotics

At the Intensive Care Unit in the University Medical Centre, Utrecht, work is going on in the development of a decision-support system that is aimed at assisting clinicians in the prescription of antibiotics to patients with ventilator-associated pneumonia (VAP) [82]. Patients admitted to an ICU are often severely ill and usually must submit to a variety of invasive medical procedures; as a consequence, these patients are generally more vulnerable to infectious diseases than healthy persons. One of the most frequently occurring infectious diseases within ICUs is pneumonia, with reported rates between 15% and 20% of all patients admitted. The decision-support system is currently implemented in the ICU, where clinicians are using a clinical information system that fully replaces the traditional paper records. The knowledge about VAP has been represented in a model-based fashion, using Bayesian networks together with decision theory as basic technology. The global structure of the model is shown in Figure 16. In particular, temporal causal knowledge has been used in developing this Bayesian network.

A preliminary laboratory evaluation of the Bayesian network and decision-theoretic model has been carried out, and yielded promising results [82]. 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

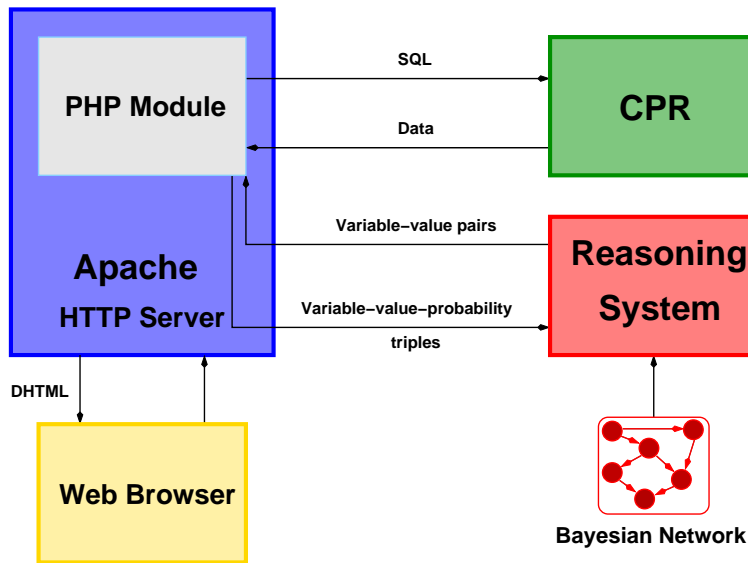


Figure 17: Architecture of the decision-support system that has been integrated with the clinical information system C2000 of Eclipsys. CPR stands for Computer-based Patient Record system, i.e. C2000.

trials.

The overall architecture of the present decision-support system is shown in Figure 17. 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 set-up 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 [131].

4 Conclusions

Even though descriptions in papers and textbooks in biology and medicine are usually qualitative in nature, at the moment only limited use is made of qualitative-reasoning technology in these fields. One reason for this may be that qualitative-reasoning technology has originally been proposed within the field of computing science, in particular artificial intelligence, and since the knowledge exchange

between computing science and fields such as medical informatics, bioinformatics, biophysics, biochemistry, biomedical statistics is not particularly strong, qualitative approaches to modelling have only slowly penetrated into the biomedical fields. In addition, qualitative methods have to compete with traditional, well-understood quantitative methods. Finally, only few researchers are familiar with this technology. There is now a clear tendency in research to combine qualitative and quantitative methods, and various examples of this have been discussed in this report. This may improve the transfer of qualitative-reasoning technology to practitioners in the Life Sciences, also in the context of model-based reasoning.

Researchers who are familiar with this technology are now increasingly exploring the use of model-based and qualitative reasoning approaches to solve actual real-world biomedical problems; some companies have also developed interest in the potential of these methods. An example is BAYESWARE¹, a firm that takes Bayesian-network technology as a basis for bioinformatics research. Some Life Science companies are actively exploring model-based and qualitative-reasoning technology as well.

References

- [1] A. Abu-Hanna and W.N.H. Jansweijer. Modeling application domain knowledge using explicit conceptualization, *IEEE-Expert* 9(5), 53–64, 1994.
- [2] A. Abu-Hanna and P.J.F. Lucas. Prognostic models in medicine: AI and statistical approaches. *Methods of Information in Medicine*, 40, 1–5, 2001.
- [3] K.-P. Adlassnig and W. Horak. Development and retrospective evaluation of HEPAXPERT-I: a routinely used expert system for interpretive analysis of hepatitis A and B serologic findings, *Artificial Intelligence in Medicine* 7, 1–24, 1995.
- [4] J.S. Aikins, J.C. Kunz, E.H. Shortliffe and R.J. Fallat. PUFF: an expert system for interpretation of pulmonary function data, in: B.C. Clancey, E.H. Shortliffe, eds. *Readings in Medical Artificial Intelligence: The First Decade* (Addison-Wesley, Reading, Massachusetts, 1984).
- [5] T.C. Aitchison, J.M. Sirel, D.C. Watt, R.M. MacKie. Prognostic trees to aid prognosis in patients with cutaneous malignant melanoma, *BMJ* 311, 1536–1539, 1995.
- [6] T. Akutsu, S. Miyano, and S. Kuhara. Inferring qualitative relations in genetic networks and metabolic pathways. *Bioinformatics*, 16(8), 727–734, 2000.
- [7] S. Andreassen, M. Woldbye, B. Falck, S.K. Andersen. MUNIN — A causal probabilistic network for interpretation of electromyographic findings, in: J. McDermott, ed., *Proceedings of the 10th International Joint Conference on Artificial Intelligence* (Morgan Kaufmann, Los Altos, CA, 1987) 366–372.
- [8] S. Andreassen. Planning of therapy and tests in causal probabilistic networks. *Artif Intell Med* 1992; 4: 227–241.
- [9] S. Andreassen, C. Riekehr, B. Kristensen, H.C. Schönheyder, L. Leibovici. Using probabilistic and decision-theoretic methods in treatment and prognosis modeling. *Artif Intell Med* 1999; 15: 121–134.

¹<http://www.bayesware.com>

- [10] C. Bailey-Kellogg, F. Zhao, and K. Yip. Spatial aggregation: Language and applications. *Proc. AAAI-96* (Morgan Kaufmann, Los Altos, 1996) 517–522.
- [11] R. Bellazzi, R. Guglielmann, and L. Ironi. A qualitative-fuzzy framework for nonlinear black-box system identification. In F. Zhao and K. Yip, editors, *Qualitative Reasoning - The Twelfth Int. Workshop*, pages 10–20, Cape Cod, MA, 1998. AAAI Press.
- [12] R. Bellazzi, R. Guglielmann, and L. Ironi. A hybrid input-output approach to model metabolic systems: An application to intracellular thiamine kinetics. *Journal of Biomedical Informatics*, 34, 221–248, 2001.
- [13] R. Bellazzi, L. Ironi, R. Guglielmann, and M. Stefanelli. Qualitative models and fuzzy systems: an integrated approach for learning from data. *Artificial Intelligence in Medicine*, 14, 5–28, 1998.
- [14] D. Berleant and B. Kuipers. Qualitative and quantitative simulation: Bridging the gap. *Artificial Intelligence*, 95, 215–256, 1997.
- [15] I. Bratko, I. Mozetič and N. Lavrač. *KARDIO: A Study in Deep and Qualitative Knowledge for Expert Systems* (The MIT Press, Cambridge, Massachusetts, 1989).
- [16] L. Breiman, J.H. Friedman, R.A. Olsen and C.J. Stone. *Classification and regression trees* (Wadsworth, Belmont, 1984).
- [17] V. Brusoni, L. Console, P. Terenziani, D. Theseider Dupré. Temporal model-based diagnosis: an overview from an abductive perspective, in: *Proceedings of CESA'96 IMACS Multiconference: Symposium on Modelling, Analysis and Simulation 1*, 326–331, 1996.
- [18] G. Carrault, M.-O. Cordier, R. Quiniou, M. Garreau, J.-J. Bellanger, and A. Bardou. A model-based approach for learning to identify cardiac arrhythmias. In W. Horn, Y. Shahar, G. Lindberg, S. Andreassen, and J. Wyatt, editors, *Proceedings of AIMDM-99 : Artificial Intelligence in Medicine and Medical Decision Making*, volume 1620 of *LNAI*, pages 165–174, Aalborg, Denmark, June 1999. Springer Verlag.
- [19] E.R. Carson, C. Cobelli, and L. Finkenstein. *The Mathematical Modeling of Metabolic and Endocrine Systems*. Wiley, New York, 1983.
- [20] G. Citro, G. Banks, G. Cooper. INKBLOT: a neurological diagnostic decision support system integrating causal and anatomical knowledge, *Artificial Intelligence in Medicine*, 10, 257–267, 1997.
- [21] P. Colli Franzone, L. Guerri, and M. Pennacchio. 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, 1998.
- [22] L. Console, D. Theseider Dupré and P. Torasso. A theory of diagnosis for incomplete causal models, in *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, 1311–1317, 1989.
- [23] L. Console, D. Theseider Dupré and P. Torasso. On the relationship between abduction and deduction, *Journal of Logic and Computation*, 1(5), 661–690, 1991.

- [24] L. Console and P. Torasso. A spectrum of logical definitions of model-based diagnosis, *Computational Intelligence* **7**(3), 133–141, 1991.
- [25] P.T. Cox and T. Pietrzykowski. General diagnosis by abductive inference, in: *Proceedings of the IEEE Symposium on Logic Programming*, 183–189, 1987.
- [26] F.T. de Dombal, D.J. Leaper, J.R. Staniland, A.P. McAnn and J.C. Horrocks. Computer-aided diagnosis of acute abdominal pain, *British Medical Journal*, **ii**, 9–13, 1972.
- [27] F.T. de Dombal, V. Dallos and W.A. McAdam. Can computer-aided teaching packages improve clinical care in patients with acute abdominal pain? *British Medical Journal* **302** (1991) 1495–1497.
- [28] P. Domingos and M. Pazzani. On the optimality of the simple Bayesian classifier under zero-one loss. *Machine Learning*, **29**, 103–130, 1997.
- [29] C. Dousson, P. Gaborit, and M. Ghallab. Situation recognition: representation and algorithms. In *Proceedings of IJCAI-93*, pp. 166–172, Chambéry, France, 1993.
- [30] C.P. Fall, E.S. Marland, J.M. Wagner, and J.J. Tyson, editors. *Computational Cell Biology*. Springer-Verlag, New York, 2002.
- [31] B. Faltings and P. Struss, editors. *Recent Advances in Qualitative Physics*. MIT Press, Cambridge, MA, 1992.
- [32] M.B. First, B.J. Weimer, S. McLinden and R.A. Miller. LOCALIZE: computer-assisted localization of peripheral nervous system lesions, *Computers and Biomedical Research*, **15**, 525–543, 1982.
- [33] K.D. Forbus. Qualitative process theory. *Artificial Intelligence*, **24**:85–168, 1984.
- [34] N.I.R. Friedman, D. Geiger and M. Goldszmidt. Bayesian network classifiers. *Machine Learning*, **29**, 131–163, 1997.
- [35] Friedman NIR, Linial M, Nachman I, Pe’er D. Using Bayesian network to analyze expression data. *J Computational Biology* 2000; **7**:601–620.
- [36] S.F. Galán, F. Aguado, F.J. Díez, J. Mira. NasoNet: joining Bayesian networks and time to model nasopharyngeal cancer spread. In: *Proceedings of the 8th Conference on Artificial Intelligence in Medicine in Europe (AIME2001)*, LNAI2101. Berlin: Springer-Verlag; 2001. p. 207–216.
- [37] G.A. Gorry and G.O. Barnett, Experience with a model of sequential diagnosis, *Computers and Biomedical Research* **1** (1968) 490–507.
- [38] F.E. Harrell, K.L. Lee, D.B. Matchar and T.A. Reichert. Regression models for prognostic prediction: advantages, problems, and suggested solutions, *Cancer Treatment Reports*, **69**(10), 1071–1077, 1985.
- [39] J. Hasty, D. McMillen, F. Isaacs, and J.J. Collins. Computational studies of gene regulatory networks: *in numero* molecular biology. *Nature Review Genetics*, **2**(4), 268–279, 2001.

- [40] <http://www-helix.inria.fr/gna>.
- [41] <http://igweb.integratogenomics.com/GOID/>.
- [42] M. Hauskrecht. Dynamic decision making in stochastic partially observable medical domains: ischemic heart disease example, in: E. Keravnou, C. Carbay, R. Baud, J. Wyatt, eds., *Proc. 6th Conf. on Artificial Intelligence in Medicine, AIME'97*, Lecture Notes in Artificial Intelligence 1211, (Springer-Verlag, Berlin, 1997) 296–299.
- [43] K.R. Heidtke and S. Schulze-Kremer. Design and implementation of a qualitative simulation model of λ phage infection. *Bioinformatics*, 14(1), 81–91, 1998.
- [44] R. Heinrich and S. Schuster. *The Regulation of Cellular Systems*. Chapman & Hall, New York, 1996.
- [45] A. Hernández, G. Carrault, F. Mora, and A. Bardou. 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, 2000.
- [46] X. Huang and F. Zhao. Relation-based aggregation: finding objects in large spatial datasets. *Intelligent Data Analysis*, 4, 129–147, 2000.
- [47] L. Ironi and M. Stefanelli. Generating explanations of pathophysiological system behaviors from qualitative simulation of compartmental models. In P. Barahona, M. Stefanelli, and J. Wyatt, editors, *Artificial Intelligence in Medicine*, volume 934 of *Lecture Notes in Computer Science*, pages 115–126. Springer-Verlag, Berlin, 1995.
- [48] L. Ironi, M. Stefanelli, and G. Lanzola. Qualitative models in medical diagnosis. *Artificial Intelligence in Medicine*, 2, 85–101, 1990.
- [49] L. Ironi and S. Tentoni. The Spatial Aggregation approach in a contouring context: problems and solutions. (submitted, 2003).
- [50] L. Ironi and S. Tentoni. On the problem of spatial adjacency relations in the Spatial Aggregation approach (in preparation).
- [51] K.J. Isselbacher, E. Braunwald, J.D. Wilson, J.B. Martin, A.S. Fauci and D.L. Kasper, eds. *Harrison's Principles of Internal Medicine* 13th edition (McGraw-Hill, New York, 1994).
- [52] F.V. Jensen. *Bayesian Networks and Decision Graphs*, Springer, New York, 2001.
- [53] H. de Jong. Modeling and simulation of genetic regulatory systems: A literature review. *Journal of Computational Biology*, 9(1), 69–105, 2002.
- [54] H. de Jong, J. Geiselman, C. Hernandez, and M. Page. Genetic Network Analyzer: Qualitative simulation of genetic regulatory networks. *Bioinformatics*, 2002. In press.
- [55] H. de Jong, J.-L. Gouzé, C. Hernandez, M. Page, T. Sari, and H. Geiselman. Dealing with discontinuities in the qualitative simulation of genetic regulatory networks. In F. van Harmelen, editor, *Proceedings of 15th European Conference on Artificial Intelligence, ECAI 2002*, pp. 412–416, Amsterdam, 2002. IOS Press.

- [56] H. de Jong, M. Page, C. Hernandez, and J. Geiselmann. Qualitative simulation of genetic regulatory networks: Method and application. In B. Nebel, editor, *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, IJCAI-01*, pages 67–73, San Mateo, CA, 2001. Morgan Kaufmann.
- [57] L. Glass and S.A. Kauffman. The logical analysis of continuous non-linear biochemical control networks. *Journal of Theoretical Biology*, 39, 103–129, 1973.
- [58] D.E. Heckerman, E.J. Horvitz and B.N. Nathwani. Towards normative expert systems: part I – The Pathfinder project. *Methods of Information in Medicine* 1992; 31: 90–105.
- [59] D.E. Heckerman and B.N. Nathwani. Towards normative expert systems: part II – probability-based representations for efficient knowledge acquisition and inference. *Methods of Information in Medicine* 1992; 31:106–116.
- [60] P.D. Karp. Hypothesis formation as design. In J. Shrager and P. Langley, editors, *Computational Models of Scientific Discovery and Theory Formation*, pp. 275–317. Morgan Kaufmann, San Mateo, CA, 1991.
- [61] P.D. Karp. Design methods for scientific hypothesis formation and their application to molecular biology. *Machine Learning*, 12, 89–116, 1993.
- [62] P.D. Karp. A qualitative biochemistry and its application to the regulation of the tryptophan operon. In L. Hunter, editor, *Artificial Intelligence and Molecular Biology*, pp. 289–324. AAAI Press/MIT Press, Menlo Park, CA/Cambridge, MA, 1993.
- [63] J. Keppens and Q. Shen. On compositional modelling. *Knowledge Engineering Review*, 16(2), 157–200, 2001.
- [64] H. Kitano, ed. *Foundations of System Biology*. MIT Press, Cambridge, MA, 2001.
- [65] J. de Kleer, A.K. Mackworth and R. Reiter. Characterizing diagnoses and systems, *Artificial Intelligence*, 52, 197–222, 1992.
- [66] W.A. Knaus, E.A. Draper, J. Lynn. Short term morbidity predictions for critically ill hospitalised patients – science and ethics, *Science*, 254, 389-94, 1991.
- [67] K. Koile and C.G. Overton. A qualitative model for gene expression. In *Proceedings of the 1989 Summer Computer Simulation Conference*, pp. 415–421. Society for Computer Simulation, 1989.
- [68] M. Korver and P.J.F. Lucas. Converting a rule-based expert system into a belief network, *Medical Informatics*, 18(3), 219–241, 1993.
- [69] B. Kuipers. Qualitative simulation. *Artificial Intelligence*, 29, 289–388, 1986.
- [70] B. Kuipers. *Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge*. MIT Press, Ca.
- [71] B.J. Kuipers. *Qualitative Reasoning* (The MIT Press, Cambridge, Massachusetts, 1994).

- [72] S.L. Lauritzen and D.J. Spiegelhalter, Local computations with probabilities on graphical structures and their application to expert systems, *Journal of the Royal Statistical Society (Series B)*, **50**, 157–224, 1987.
- [73] D.J. Lockhart and E.A. Winzeler. Genomics, gene expression and DNA arrays. *Nature*, 405, 827–836, 2000.
- [74] W.S. Lovejoy. A survey of algorithmic methods for partially observable Markov decision processes, *Annals of Operations Research*, 28, 47–66, 1991.
- [75] P.J.F. Lucas. Logic engineering in medicine, *The Knowledge Engineering Review*, 10 (2), 153–179, 1995.
- [76] P.J.F. Lucas, H. Boot and B. Taal. A decision-theoretic network approach to treatment management and prognosis, in: A. Macintosh and R. Milne, eds., *Applications and Innovations in Expert Systems V* (SGE
- [77] P.J.F. Lucas, H. Boot H, B.G. Taal. Computer-based decision-support in the management of primary gastric non-Hodgkin lymphoma. *Methods of Information in Medicine* 1998;37:206–219. S, Cambridge, 1997) 17–30.
- [78] P.J.F. Lucas. The representation of medical reasoning models in resolution-based theorem provers, *Artificial Intelligence in Medicine*, **5**, 395–414, 1993.
- [79] P.J.F. Lucas. Analysis of notions of diagnosis, *Artificial Intelligence*, **105**(1-2), 293–341, 1998.
- [80] P.J.F. Lucas. Symbolic diagnosis and its formalisation, *The Knowledge Engineering Review*, **12**(2), 109–146, 1997.
- [81] P.J.F. Lucas, A.M. Tholen and G. van Oort. An intelligent system for pacemaker reprogramming. *Artificial Intelligence in Medicine*, **17**, 249–269, 1999.
- [82] 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.
- [83] P.J.F. Lucas, R. Kuipers, F. Feith. A system for pacemaker treatment advice. In: R. López de Mántaras and L. Saitta (Eds.), *Proceedings of the 16th ECAI 2004*, IOS Press, Amsterdam, 2004, pp. 735–739.
- [84] H.H. McAdams and A. Arkin. Simulation of prokaryotic genetic circuits. *Annual Review of Biophysics and Biomolecular Structure*, 27:199–224, 1998.
- [85] H.H. McAdams and L. Shapiro. Circuit simulation of genetic networks. *Science*, 269:650–656, 1995.
- [86] T. Mestl, E. Plahte, and S.W. Omholt. A mathematical framework for describing and analysing gene regulatory networks. *Journal of Theoretical Biology*, 176:291–300, 1995.
- [87] R.A. Miller, H.E. Pople and J.D. Myers. INTERNIST-1, an experimental computer-based diagnostic consultant for general internal medicine, *New England Journal of Medicine* **307** (1982) 468–476.

- [88] G.B. Moody and R.G. Mark. The mit-bih arrhythmia database on cd-rom and software for use with it. In *Proceedings of Computers in Cardiology Conference*, volume 17, pages 185–188. IEEE Computer Society Press, 1990.
- [89] I. Mozetič. Hierarchical model-based diagnosis, in: W. Hamscher, L. Console and J. de Kleer, eds., *Readings in Model-based Diagnosis* (Morgan Kaufmann, San Mateo, 1992) 354–372.
- [90] S. Muggleton. Inductive logic programming: issues, results and the challenge of learning language in logic. *Artificial Intelligence*, 114(1-2):283–296, 1999.
- [91] J. Neuberger. Prognostic models – what is their future?, *European Journal of Gastroenterology & Hepatology* 9 (1997) 1145–1147.
- [92] A. Pandey and M. Mann. Proteomics to study genes and genomes. *Nature*, 405:837–846, 2000.
- [93] R.S. Patil, P. Szolovits and W.B. Schwartz. Modeling knowledge of the patient in acid-base and electrolyte disorders, in: P. Szolovits, ed., *Artificial Intelligence in Medicine* (Westview Press, Boulder, Colorado, 1982) 191–226.
- [94] Z. Pawlak, *Rough Sets - Theoretical Aspects of Reasoning about Data* (Kluwer Academic Publishers, Dordrecht, 1991).
- [95] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: networks of plausible inference* (Morgan Kaufmann, San Mateo, CA, 1988).
- [96] Y. Peng and J.A. Reggia. *Abductive Inference Models for Diagnostic Problem Solving* (Springer-Verlag, New York, 1990).
- [97] D. Poole, R. Goebel and R. Aleliunas. Theorist: a logical reasoning system for defaults and diagnosis, in: N. Cercone and G. Mc Calla, eds., *The Knowledge Frontier* (Springer-Verlag, Berlin, 1987) 331–352.
- [98] H.E. Pople. Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnosis, in: P. Szolovits, ed., *Artificial Intelligence in Medicine* (Westview Press, Boulder Colorado, 1982) 119–190.
- [99] M.L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley, New York, 1994).
- [100] R. Quiniou, M.-O. Cordier, G. Carrault, and F. Wang. Application of ILP to cardiac arrhythmia characterization for chronicle recognition. In C. Rouveirol and M. Sebag, editors, *Proceedings of ILP-2001*, volume 2157 of *LNAI*, pages 193–215. Springer Verlag, september 2001.
- [101] L. De Raedt and W. Van Laer. Inductive constraint logic. In *Proceedings of ALT-95 (6th Workshop on Algorithmic Learning Theory)*, volume 997 of *LNAI*, pp. 80–94, 1995.
- [102] H. Raiffa. *Decision Analysis: Introductory Lectures on Choices under Uncertainty* (Addison-Wesley, Reading, MA, 1968).
- [103] M. Ramoni, P. Sebastiani, P. Cohen. Bayesian clustering by dynamics. *Machine Learning* 2002; 47: 91–121.

- [104] J.A. Reggia, D.S. Nau and Y. Wang. Diagnostic expert systems based on a set covering model, *International Journal of Man Machine Studies*, **19**, 437–460, 1983.
- [105] R. Reiter. A theory of diagnosis from first principles, *Artificial Intelligence*, **32**, 57–95, 1987.
- [106] M. Santillán and M.C. Mackey. Dynamic regulation of the tryptophan operon: A modeling study and comparison with experimental data. *Proceedings of the National Academy of Sciences of the USA*, 98(4), 1364–1369, 2001.
- [107] E. Santos Jr and E.S. Santos. Polynomial solvability of cost-based abduction, *Artificial Intelligence*, **86**, 157–170, 1996.
- [108] R.D. Shachter. Evaluating influence diagrams, *Operations Research*, 34(6), 871–882, 1986.
- [109] E.H. Shortliffe. *Computer-based Medical Consultations: MYCIN* (Elsevier, New York, 1976).
- [110] E.H. Shortliffe. Knowledge-based systems in medicine, in: K.-P. Adlassnig, G. Grabner, S. Bengtsson and R. Hansen, eds., *Proceedings of Medical Informatics Europe 1991*, Lecture Notes in Medical Informatics **45** (Springer-Verlag, Berlin, 1991) 5–9.
- [111] P. Smolen, D.A. Baxter, and J.H. Byrne. Modeling transcriptional control in gene networks: Methods, recent results, and future directions. *Bulletin of Mathematical Biology*, 62, 247–292, 2000.
- [112] E.H. Snoussi. Qualitative dynamics of piecewise-linear differential equations: A discrete mapping approach. *Dynamics and Stability of Systems*, 4(3-4), 189–207, 1989.
- [113] H.C. Sox, M.A. Blatt, M.C. Higgins, K.I. Marton. *Medical Decision Making* (Butterworths, Boston, 1988).
- [114] D.J. Spiegelhalter and R.P. Knill-Jones. Statistical and knowledge-based approaches to clinical decision-support systems, with an application in gastroenterology, *Journal of the Royal Statistical Society*, **147**, 35–77, 1984.
- [115] W.W. Steyerberg, H.J. Keizer, S.D. Fossa, et al. Prediction of residual retroperitoneal mass histology following chemotherapy for metastatic nonseminomatous germ cell tumor: multivariate analysis of individual patient data from 6 study groups, *Journal of Clinical Oncology*, 13, 1177–1187, 1995.
- [116] B. Taccardi, B.B. Punske, R.L. Lux, R.S. MacLeod, P.R. Ershler, T.J. Dustman, and Y. Vyhmeister. Useful lessons from body surface mapping. *Journal of Cardiovascular Electrophysiology*, 9(7), 773–786, 1998.
- [117] A. ten Teije and F. van Harmelen. An extended spectrum of logical definitions for diagnostic systems, in: G.M. Provan, ed., *DX-94, 5th International Workshop on Principles of Diagnosis*, 334–342, 1994.
- [118] A. Theodossi, D.J. Spiegelhalter and B. Portmann et al. The value of clinical, biochemical, ultrasound and liver biopsy data in assessing patients with liver disease, *Liver*, **3**, 315–326, 1983.
- [119] R. Thomas and R. d’Ari. *Biological Feedback*. CRC Press, Boca Raton, FL, 1990.

- [120] P. Torasso and L. Console. *Diagnostic problem solving* (North Oxford Academic Publishers, London, 1989).
- [121] R.B. Trelease, R.A. Henderson, and J.B. Park. A qualitative process system for modeling NF- κ B and AP-1 gene regulation in immune cell biology research. *Artificial Intelligence in Medicine*, 17, 303–321, 1999.
- [122] R.B. Trelease and J. Park. Qualitative process modeling of cell-cell-pathogen interactions in the immune system. *Computer Methods and Programs in Biomedicine*, 51(3), 171–181, 1996.
- [123] S. Tuhim, J. Reggia and S. Goodall. An experimental study of criteria for hypothesis plausibility, *Journal of Experimental and Theoretical Artificial Intelligence*, 3, 129–144, 1991.
- [124] Tutorial: How to Perform a Decision analysis, *Medical Decision Making*, 17, 123–159, 1997.
- [125] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*, 2nd ed. (Princeton University Press, Pinceton, NJ, 1947).
- [126] F. Wang, R. Quiniou, G. Carrault, and M.-O. Cordier. Learning structural knowledge from the ECG. In *Proceedings of ISMDA-2001*, volume 1933 of *LNCS*, pages 288–294. Springer Verlag, october 2001.
- [127] S.M. Weiss, C.A. Kulikowski, S. Amarel and A. Safir. A model-based method for computer-aided medical decision making, *Artificial Intelligence*, 11, 145–172, 1978.
- [128] D.S. Weld and J. de Kleer, editors. *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufmann, San Mateo, CA, 1990.
- [129] T.D. Wu. A problem decomposition method for efficient diagnosis and interpretation of multiple disorders, *Computer Methods and Programs in Biomedicine*, 35, 239–250, 1991.
- [130] J.C. Wyatt and D.G. Altman. Commentary – prognostic models: clinically useful or quickly forgotten, *BMJ*, 311, 1539–1541, 1995.
- [131] C.P. Friedman and J.C. Wyatt. *Evaluation Methods in Medical Informatics*. Springer-Verlag, New York, 1997.