

Prognostic Methods in Medicine

Peter Lucas

Department of Computer Science, Utrecht University,
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
E-mail: lucas@cs.uu.nl

Ameen Abu-Hanna

Department of Medical Informatics, AMC-University of Amsterdam,
Meibergdreef 15, 1105 AZ Amsterdam, The Netherlands
E-mail: A.Abu-Hanna@amc.uva.nl

Abstract

Prognosis – the prediction of the course and outcome of disease processes – plays an important role in patient management tasks like diagnosis and treatment planning. Prognostic models form, therefore, an integral part of a number of systems supporting these tasks. Furthermore, prognostic models constitute instruments to evaluate the quality of health care and the consequences of health care policies by comparing predictions according to care norms with actual results. Approaches to developing prognostic models vary from using traditional probabilistic techniques, originating from the field of Statistics, to more qualitative and model-based techniques, originating from the field of Artificial Intelligence. In this paper, various approaches to constructing prognostic models, with emphasis on methods from the field of Artificial Intelligence, are described and compared.

Keywords: medical prognostic models; medical AI; machine learning; medical decision making.

1 Role of prediction and prognosis in medicine

Scientific progress in Medicine is made at an annually increasing pace through the continuous improvement of techniques for the diagnosis and treatment of disease. Modern patient management requires that considerations of the appropriateness of medical interventions are based on scientific evidence, otherwise, the level of health care is at risk and there is little justification for the interventions and their associated costs. These principles are nowadays referred to as *Evidence-Based Medicine* (EBM).

Ideas concerning EBM are gradually starting to influence patient care: it is increasingly recognised that diagnosis and treatment of disease should not be taken as isolated events, but, instead, should be established with respect to expected outcomes both for the patient and the community as a whole. For example, establishing a diagnosis of concomitant disease in a terminally ill patient may have little or no value; the best choice here may be to leave it open. Balancing the costs and benefits of a major surgical intervention like cardiopulmonary transplantation is another example. Here, quality of life, risks of the surgical procedure, unwanted side-effects of immunosuppressive drugs, life expectancy, recurrence of disease, financial costs of the procedure and budgetary limitations should all be taken into account.

In arguing about the tradeoffs 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 [17].

In this editorial, the place of prognostic models in the medical management – diagnosis, treatment selection and follow-up – of patients with disease is sketched, and important techniques used in the construction of prognostic models are reviewed and brought in perspective. One of the main problems from the viewpoint of the AI practitioner concerns the suitability of a particular technique for building a prognostic model. Some thoughts are devoted to this important topic in the various sections of this paper, but the contributions to this thematic issue on prognostic methods in medicine, which appear in this special issue, provide valuable insights based on results of actual research projects.

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.

2.1 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 [27]. 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.

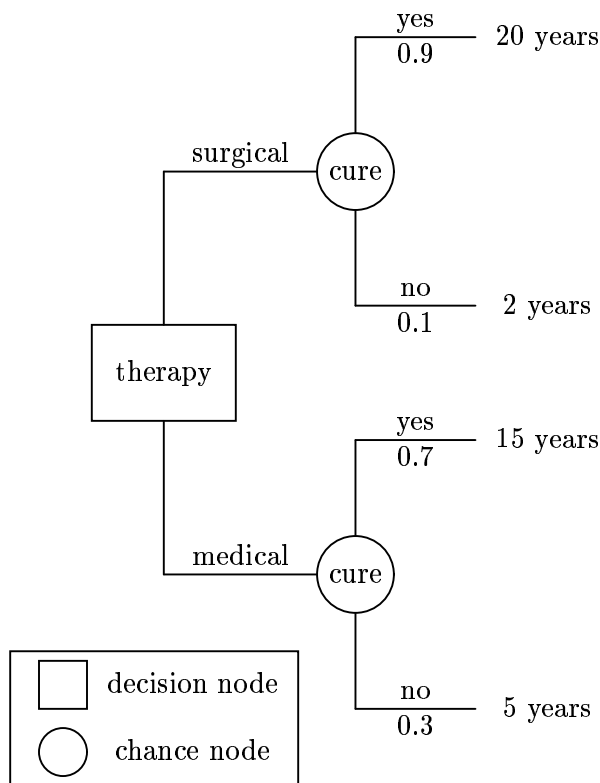


Figure 1: A simple decision tree.

Decision Trees The most popular tool to express sequential decision making is the *decision tree*¹ (see the first four parts of the tutorial in [26]). 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 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 1 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 $P(\text{cure} = \text{yes} | \text{therapy} = \text{surgical})$. The leaves of the decision tree are labeled 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) [23, 21].

¹The term decision tree in Decision Analysis does not refer to the same thing as in Machine Learning as we will see.

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 1 incorporates the probability $P(\text{cure} = \text{yes} | \text{therapy} = \text{surgical})$, but if we would like to know the probability $P(\text{therapy} = \text{surgical} | \text{cure} = \text{yes})$ (which percentage of patients who have been cured, have had surgery), the probabilities $P(\text{cure} = \text{yes})$ and $P(\text{therapy} = \text{surgical})$ would be required, because, according to Bayes' rule,

$$P(\text{therapy} = \text{surgical} | \text{cure} = \text{yes}) = \frac{P(\text{cure} = \text{yes} | \text{therapy} = \text{surgical})P(\text{therapy} = \text{surgical})}{P(\text{cure} = \text{yes})}$$

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 [9, 25]. For example, the APACHE III [11] 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 [24].

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) [20]. 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) [14]. 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




Figure 2: A simple Markov process model.

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 [10].

There are various algorithms in use in connection with Markov processes, varying from exact algebraic algorithms and dynamic programming algorithms to stochastic simulation algorithms [20].

As an example of a Markov process consider Figure 2. 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 aging on mortality. 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 Artificial intelligence 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.

2.2.1 Explicit representations

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 [15]. Anatomical

knowledge can be very useful to localise causes of malfunction, in particular in neurology [8]. 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 [6]. 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 traditionally 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* [12]. 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.

2.2.2 Model construction

Unlike traditional statistical methods where the (mathematical) structure of models is given and probabilistic information is learnt from data, a characteristic of some AI 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. Below is given an overview of some important, mainly inductive, techniques, that is, techniques for forming generic models based on given examples. Note that the term datamining, the heart of the process of Knowledge Discovery in Databases which relies heavily on inductive ML techniques is sometimes used as a synonym to inductive ML.

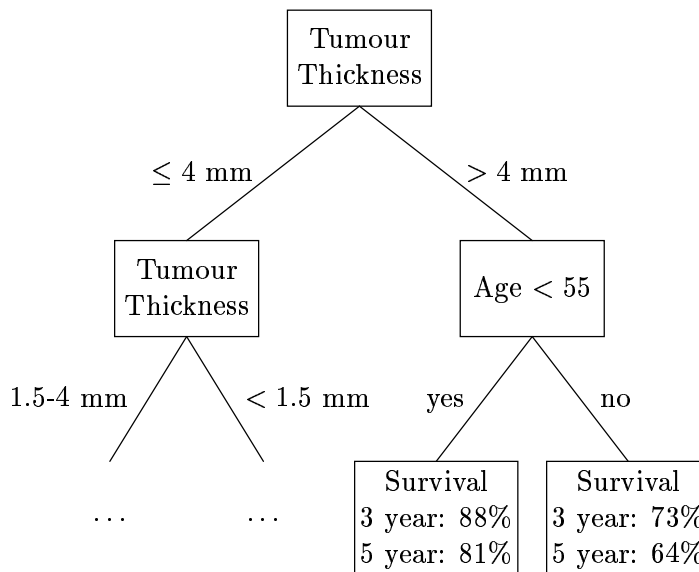


Figure 3: An ML decision tree for predicting survival based on tumour thickness and age.

2.2.3 AI techniques

Decision Trees Decision trees² (DTs) are very popular in ML applications and are also used as prognostic models in Medicine [4]. They are attractive because they provide a symbolic representation that lends itself to easy interpretation by humans. The representation can also be extended or easily modified when a tree is translated into convenient if-then rules.

To illustrate, the prognostic decision tree in Figure 3 predicts patient survival based on tumour thickness and age. The rightmost path in the tree can be represented as:

IF (*TumourThickness* > 4mm) \wedge \neg (age < 55) THEN
 (survival_3_years = 73%) \wedge (survival_5_years = 64%)

Although learning a decision tree that is ‘best’ in the sense of compactness is computationally intractable, learning a decision tree that is adequate for the purpose at hand can be done in polynomial time when special heuristics are used. The idea in DT learning is to find those attributes in the dataset that contribute most to identifying homogeneous groups (with respect to their classification of survival) in the dataset, by minimising entropy. Independently, work in statistics has led to similar techniques for inducing DTs [7].

Furthermore, decision trees do not impose a bias on the *expressiveness* of the models. This is in contrast to some statistical methods which restrict the form of the model (e.g. preassuming a linear model) or impose assumptions on the probability distributions or the (in)dependency structure among variables. Of course, when these assumptions do hold then restricting the model expressiveness does pay and in fact it lowers the danger of overfitting the data to which decision trees are sensitive. Overfitting results in a (too) good fit of the training data on the expense of prediction ability on unseen examples. Overfitting is hence an important issue in DT research and methods like pruning of nodes based on test data for validation is often used.

²Decision trees in ML are different from those in Decision Analysis as described below.

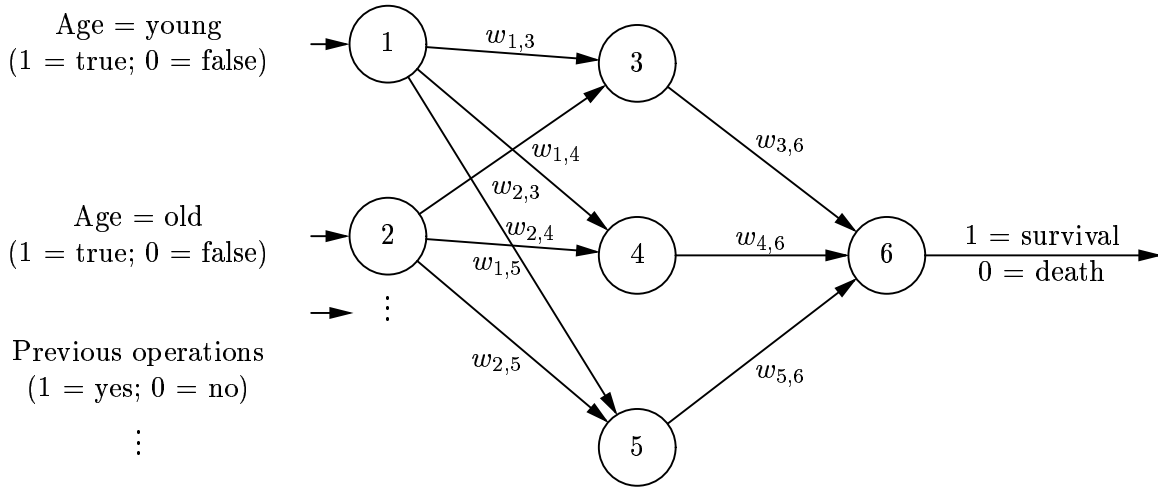


Figure 4: A 3-layer neural network.

Neural Networks A Neural network (NN) [5] is a computational representation that takes as input a vector of (e.g. real) numbers, say encoded patient’s features, and outputs a number (or a numerical vector) that is interpreted as, say, survival probability of that patient (see Figure 4).

The inputs are weighted and communicated to the nodes by edges. The nodes usually check if the sum of their weighted inputs is greater than a specific threshold and, based on the result, output 1 or 0 which are in turn weighted and conveyed to other nodes. Learning in NNs comprises of calculating the values of edge weights $w_{i,j}$ and node thresholds that can best explain the data. The learning algorithms are based on manipulating the weights and thresholds so as to minimise the error resulting from the discrepancy between the predicted value of the outcome according to the model and its expected value as indicated by a training set.

NNs are robust to noise in the training set and are capable of expressing complicated interactions, including nonlinear ones, among variables and hence are more flexible than regression analysis in statistics, where the particular type of function to be learnt (for example an exponential function $f(x) = ae^{bx}$) must be known beforehand. As with decision trees, one should beware of overfitting. Although quite popular in AI research, NNs are not very suitable when human understanding of the relationship captured by the NN is important. This ‘black-box’ feature makes NN less attractive in prognostic methods in Medicine, where clinical credibility is essential. However, there are techniques that translate NN representations into symbolic ones.

Bayesian Belief Networks A Bayesian belief network (BBN), also called causal probabilistic network (CPN), 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 $P(V|\pi(V))$, 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 BBN mirrors the dependencies and independencies among the variables corresponding to the nodes in the graph. For example, a variable is dependent of its parents and children, but the

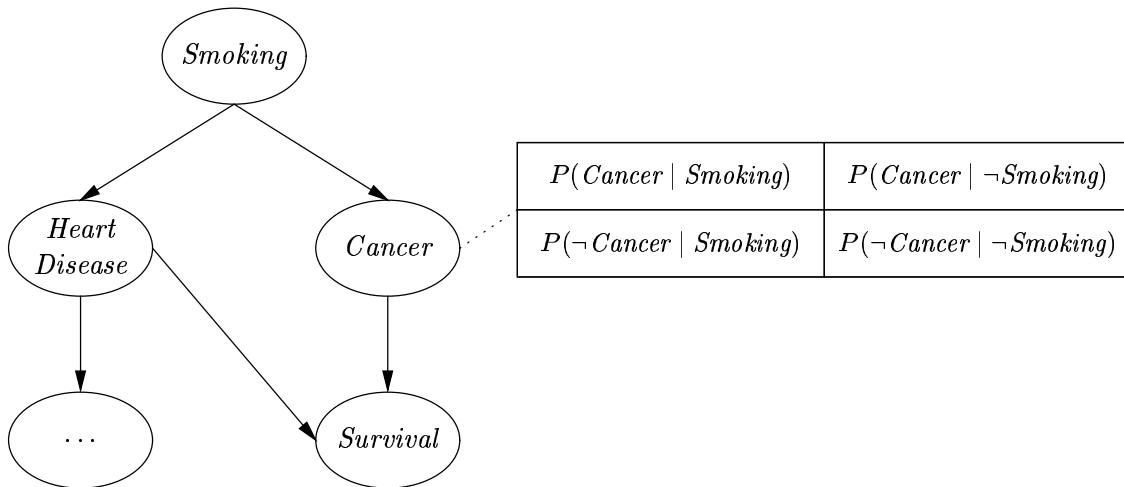


Figure 5: A Bayesian belief network.

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’. In Figure 5, for example, ‘Cancer’ is conditionally independent of ‘Heart Disease’, given the value of its immediate parent node (‘Smoking’).

BBNs are increasingly used in clinical and medical research [3, 16] and several probabilistic inference algorithms are available to process evidence, by manipulating the probabilities incorporated in a BBN [13, 19]. It is also possible to explicitly distinguish between arcs having a *temporal* meaning, expressing a transition in time, and atemporal arcs, expressing probabilistic dependencies within the same time slice. The possibility to express such temporal relationships is a useful feature for the development of prognostic models. By adding decision nodes and a value node to the chance nodes of BBNs, a formalism is obtained that can be used to determine an optimal sequence of decisions, based on maximal expected utility calculations. This formalism is known as the *influence diagram* formalism [22]. Diagnostic and prognostic information in an influence diagram are represented by BBN parts.

Genetic Algorithms Genetic Algorithms (GAs) are algorithms inspired by the theory of evolution. A new collection (generation) of hypotheses (candidate models that are to explain the data) is created by applying ‘evolutionary’ operations like random mutation and crossover on the most recent (best so far) generation of hypotheses. A hypothesis is usually described by a string of bits or symbols. For example, the prognostic hypothesis H_1 , in Figure 6 represented by the bit string ‘1 01 010’, is interpreted as:

IF $(age > 65) \wedge$
 $(1.5mm < TumourThickness) \wedge$
 $(TumourThickness \leq 2.4mm)$
 THEN $LifeExpectancy = 7years$

For each hypothesis, a numerical measure, called *fitness*, is assigned which depicts how well the hypothesis is able to explain the (training) data. A GA operates on an initial, randomly generated, pool of hypotheses. The fittest hypotheses are found and are used as the basis for the new generation of hypotheses by evolutionary operators such as random mutation (e.g.

	Age > 65	Tumour thickness		Life expectancy
H_1 :	1	01	...	010
	0: false	00: ≤ 1.5 mm		000: 1 year
	1: true	01: (1.5, 2.4] mm		\vdots
		10: (2.4, 3.6] mm		010: 7 years
		11: > 3.6 mm		\vdots

Figure 6: Hypothesis representation and interpretation in GA.

patient	age	tumour-thickness	...	survival
p_1	45	3	...	Yes
p_2	22	5	...	Yes
p_3	45	3	...	Yes
p_4	45	3	...	Yes
p_5	22	5	...	Yes
p_6	35	4	...	No
p_7	22	5	...	No
p_8	35	4	...	No

Table 1: A dataset of patient attributes (condition attributes and a decision attribute (survival)).

H_1 in Figure 6 might become ‘1 10 010’ effectively changing the conditions in the rule) and crossover between pairs of hypotheses (e.g. part of the conditions of H_1 might be replaced by those of another hypothesis). This process is iterated until satisfactorily fit hypotheses are discovered. The GA approach is attractive due to the intuitive hypothesis interpretation and the insight into the gradual improvement of the hypotheses.

Rough Set Theory Rough Set Theory [18] provides a mathematical tool for representing and reasoning about vagueness and uncertainty. The notion of *indiscernibility* plays an important role in the theory. For example, patients p_1 and p_3 in Table 1 are indiscernible with respect to *age* and *tumour-thickness* since they have the same values for these attributes.

In the rough set approach, one can approximate a set of interest in terms of a subset of attributes by means of an upper and a lower bound. For example, consider the set X of patients in Table 1 who survived, $X = \{p_1, \dots, p_5\}$, and suppose we are to use the attributes *age* and *tumour-thickness* in the approximation of X . The sets $\{p_1, p_3, p_4\}$, $\{p_2, p_5, p_7\}$ and $\{p_6, p_8\}$ consist of indiscernible elements in the entire dataset. The *B-lower approximation* of X , denoted by $\underline{B}X$, where $B = \{age, tumour-thickness\}$, is equal to the union of every indiscernible set contained in X . In our example, only the indiscernible set $\{p_1, p_3, p_4\}$ meets this criterion. The *B-upper approximation* of X , denoted by $\overline{B}X$, consists of the union of all indiscernible sets with a nonempty intersection with X with respect to the values of the attributes in B ; in this case we have that $\overline{B}X = \{p_1, p_2, p_3, p_4, p_5, p_7\}$. $\underline{B}X$ includes all elements in X for which B is sufficient for their classification (here the survival) and $\overline{B}X - \underline{B}X$, called the *B-boundary region*, comprises elements for which B is insufficient for conclusive classification.

When the dataset depicts prognostic information, rough set theory offers the tools for the

estimation of the prognostic power of various subsets of attributes – measured for example in terms of the cardinality of the B -boundary region. This approach can be used, for example, to identify attributes which correspond to costly tests which do not enhance the prognostic power achieved by other, less costly, attributes.

3 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.

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 [28]. 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.

Without any doubt, Evidence-Based Medicine will further grow in significance, and, as a consequence, the role of prognostic models and of systems that aid in the prediction of prognosis may be expected to increase in the near future.

References

- [1] A. Abu-Hanna and W.N.H. Jansweijer, Modeling application domain knowledge using explicit conceptualization, *IEEE-Expert* 9(5) (1994) 53–64.
- [2] A. Abu-Hanna and P.J.F. Lucas, Intelligent prognostic methods in medicine, in: P. Borne, M. Ksouri and A. el Kamel, eds., *Proceedings of Computational Engineering in Systems Applications 1998* (IMACS-IEEE, UCIS, Lille) 312–317.
- [3] 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.
- [4] 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 (1995) 1536–1539.
- [5] C.M. Bishop, *Neural Networks for Pattern Recognition* (Oxford University Press, Oxford, 1995).
- [6] I. Bratko, I. Mozetič and N. Lavrač, *KARDIO: A Study in Deep and Qualitative Knowledge for Expert Systems* (The MIT Press, Cambridge, Massachusetts, 1989).
- [7] L. Breiman, J.H. Friedman, R.A. Olsen and C.J. Stone, *Classification and regression trees* (Wadsworth, Belmont, 1984).
- [8] G. Citro, G. Banks, G. Cooper, INKBLOT: a neurological diagnostic decision support system integrating causal and anatomical knowledge, *Artificial Intelligence in Medicine* 10 (1997) 257–267.
- [9] 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) (1985) 1071–1077.
- [10] 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.
- [11] W.A. Knaus, E.A. Draper, J. Lynn, Short term morbidity predictions for critically ill hospitalised patients – science and ethics, *Science* 254 (1991) 389–94.
- [12] B.J. Kuipers, *Qualitative Reasoning* (The MIT Press, Cambridge, Massachusetts, 1994).
- [13] S.L. Lauritzen, 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 (1987) 157–224.
- [14] W.S. Lovejoy, A survey of algorithmic methods for partially observable Markov decision processes, *Annals of Operations Research* 28 (1991) 47–66.

- [15] P.J.F. Lucas, Logic engineering in medicine, *The Knowledge Engineering Review* 10 (2) (1995) 153–179.
- [16] 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* (SGES, Cambridge, 1997) 17–30.
- [17] J. Neuberger, Prognostic models – what is their future?, *European Journal of Gastroenterology & Hepatology* 9 (1997) 1145–1147.
- [18] Z. Pawlak, *Rough Sets - Theoretical Aspects of Reasoning about Data* (Kluwer Academic Publishers, Dordrecht, 1991).
- [19] J. Pearl, *Probabilistic Reasoning in Intelligent Systems* (Morgan Kaufman, San Mateo, California, 1988).
- [20] M.L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley, New York, 1994).
- [21] H. Raiffa, *Decision Analysis: Introductory Lectures on Choices under Uncertainty* (Addison-Wesley, Reading, MA, 1968).
- [22] R.D. Shachter, Evaluating influence diagrams, *Operations Research* 34 (6) (1986) 871–882.
- [23] H.C. Sox, M.A. Blatt, M.C. Higgins, K.I. Marton, *Medical Decision Making* (Butterworths, Boston, 1988).
- [24] D.J. Spiegelhalter, 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 (1984) 35–77.
- [25] 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 (1995) 1177–1187.
- [26] Tutorial: How to Perform a Decision analysis, *Medical Decision Making* 17 (1997) 123–159.
- [27] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, 2nd ed. (Princeton University Press, Princeton, NJ, 1947).
- [28] J.C. Wyatt and D.G. Altman, Commentary – prognostic models: clinically useful or quickly forgotten, *BMJ* 311 (1995) 1539–1541.