Bayesian Networks in Medicine: a Model-based Approach to Medical Decision Making

Peter Lucas
Department of Computing Science
University of Aberdeen
Scotland, UK
plucas@csd.abdn.ac.uk

Abstract

Bayesian networks have been introduced in the 1980s. Research to explore the use of the formalism in the context of medical decision making started in the 1990s. The formalism possesses the unique quality of being both a statistical and an AI-like knowledge-representation formalism. As it allows for structuring domain knowledge, by exploiting causal and other relationships between domain variables, the formalism is also model-based. In this paper the use of the formalism in building medical decision support systems in medicine is discussed, taking the problem of optimal prescription of antibiotics to patients with pneumonia in the ICU as a real-life example.

Keywords: medical decision support, intelligent systems, Bayesian networks.

1 Introduction

Bayesian networks (BNs) have been introduced in the 1980s as a formalism for representing and reasoning with models of problems involving uncertainty, adopting probability theory as a basic framework [7]. Since the beginning of the 1990s researchers are exploring its possibilities for developing medical applications.

The BN formalism offers a natural way to represent the uncertainties involved in medicine when dealing with diagnosis, treatment selection, planning, and prediction of prognosis [3]. This is due to the fact that the influences and probabilistic interactions among variables can be described readily in a BN. As the formalism is declarative in nature, any (often conditional) probabilistic statement can be computed from a given BN, where the statement may concern both individual and combinations of variables. This allows asking questions such as "What is likely to be the result for the patient if I decide to request this test, to prescribe this treatment and so on". Another attractive feature of the formalism is that it is closely related to causal models, which explains why some researchers refer to it as the *causal probabilistic network (CPN) formalism*. An actual BN can often be understood in terms of cause-effect relationships.

In this paper, BNs are discussed from the viewpoint of their use in medical decision making, in particular diagnosis, (prognostic) prediction and treatment selection. A BN model that was developed to assist clinicians in the diagnosis and selection of antibiotic treatment for patients with pneumonia in the ICU is taken as a running example [5].

2 Modelling

Developing a model of a realistic medical problem is usually far from easy, and using Bayesian networks yields no exception in this respect. As is the case with other representation formalisms, there are particular guidelines which facilitate developing a BN [4]. We start by summarising some facts concerning the running example of this paper.

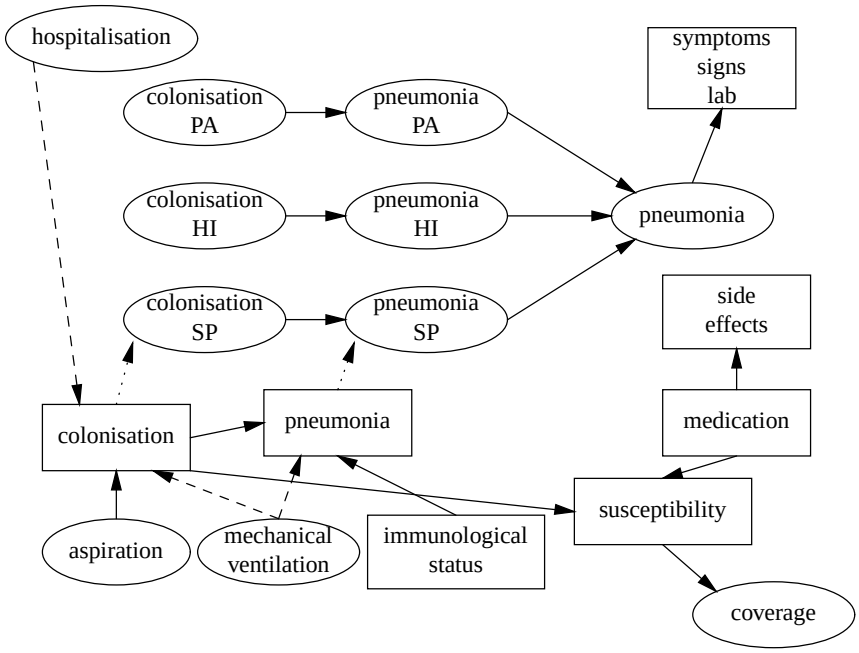

Figure 1: Detailed structure of part of the VAP model. Only three of the microorganisms included in the model are shown. Boxes stand for collections of similar vertices. Dotted arcs point to the actual topology of the network. Solid arcs stand for atemporal stochastic influences, whereas dashed arcs indicate temporal influences. Abbreviations of names of bacteria: PA = Pseudomonas aeruginosa, HI = Haemophilus influenzae, SP = Streptococcus pneumoniae.

2.1 Ventilator-associated pneumonia

Many patients admitted to an ICU need respiratory support by mechanical ventilation; furthermore, many of these patients are severely ill, which may negatively affect the functionality of their immune system. Both conditions promote the development of bacterial pneumonia [1]. Pneumonia is a very common disease in ICU patients; especially, ventilator-associated pneumonia (VAP) may arise in patients who are mechanically ventilated. Because of the wide-spread dissemination of multi-resistant bacteria in hospitals and ICUs in particular, with which patients will become colonised after some time, effective treatment of VAP is seen as an issue of major concern. Choosing the 'right' therapy means selecting antibiotics that are effective against the causative organisms, without causing major side effects.

2.2 The Bayesian-network model

Figure 1 gives an overview of the structure of the BN model of VAP which was developed together with infectious disease specialists. The structure of a BN can be designed using knowledge of known causal dependences, influences or correlations. All or part of these may be derived from knowledge of domain experts, obtained from descriptions in literature, or extracted from data using structure-learning algorithms. Formally, a Bayesian network $\mathcal{B} = (G, \Pr)$ is a directed acyclic graph $G = (V(G), A(G))$ with set of vertices $V(G) = \{V_1, \ldots, V_n\}$, representing stochastic variables, and a set of arcs $A(G) \subseteq V(G) \times V(G)$, representing stochastic dependences and independences among the variables. On the set of stochastic variables a joint probability distribution $\Pr(V_1, \ldots, V_n)$ is defined that respects the (in)dependences represented in the graph: $\Pr(V_1, \ldots, V_n) = \prod_{i=1}^{n} \Pr(V_i \mid \pi(V_i))$, where $\pi(V_i)$ stands for the variables corresponding to the parents of vertex V_i. One of the attractive features of BNs is that it is possible to combine information from various sources, for example starting by defining a probability distribution from one source, and then refining it using data.

An important role in the model is played by the temporal process of *colonisation* of the airways by pathogens. The fact that this process is temporal, is expressed by the interaction between duration of stay

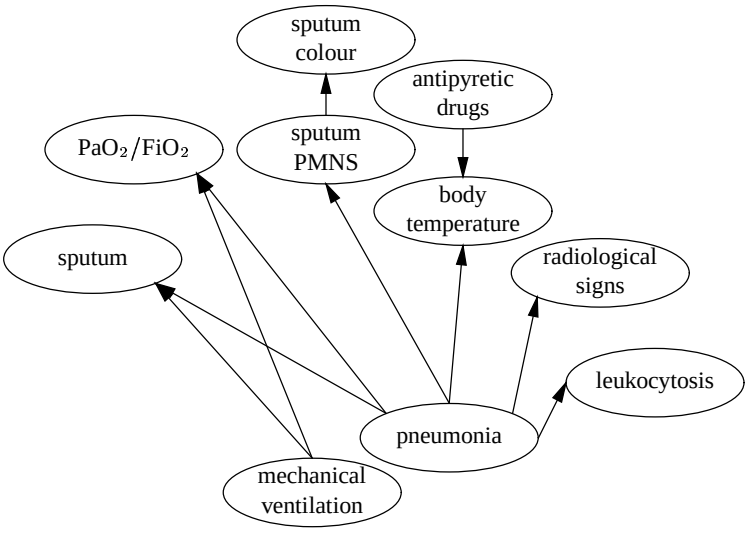

Figure 2: Probabilistic model of signs and symptoms of pneumonia.

(hospitalisation) and duration of mechanical ventilation: both are positively correlated to colonisation with pathogens. It is in principle also possible to represent temporal knowledge by means of temporal arcs between the same variables at different points in time, but reasoning with such a representation, which resembles a Markov process, may be very demanding computationally. Other arcs have a causal reading without a strong temporal connotation. For example, aspiration of stomach content is another factor positively correlated to colonisation with some pathogens (*Enterobacteriaceae*). When a patient gets colonised with a pathogen, there is a certain probability that pneumonia will develop caused by that specific pathogen. Therefore, an arc is drawn from 'colonisation' to 'pneumonia'. Duration of mechanical ventilation and the immunological status of a patient influence the probability that pneumonia will arise as well; therefore, an arc is drawn from 'immunological status' and 'mechanical ventilation' to 'pneumonia'. When a patient is affected by pneumonia, symptoms and signs can be observed, as well as abnormalities in laboratory values; this part of the model is shown in Figure 2. Here the arcs sometimes have a causal reading and sometimes the less specific meaning of a correlational influence.

Graphs like the one shown in Figure 1 appear easy to understand, but their underlying formal semantics is sophisticated. For example, the structure in Figure 2 tells us that leucocytosis is conditionally independent of body temperature given presence or absence of pneumonia. The notion of *induced* dependence is also central to the theory; it signifies a (dynamic) change in the dependence relation represented by the graph. For example, the various colonisation variables are (unconditionally) independent, but will become dependent once information on the presence or absence of their common consequence, pneumonia, is entered into the network. Insufficient understanding of the formal meaning of BNs may give rise to modelling flaws.

Prescription of antibiotic therapy amounts to selecting none, one or two antibiotic drugs, modelled by two identical therapy vertices. Let d be the number of possible drugs (including none). Note that with two treatment vertices d^2 combinations are possible, of which $\binom{d-1}{2}$ are unique; the total number of different (bi- and mono-)therapies that actually can be prescribed is thus $\binom{d-1}{2} + d$.

3 Medical problem solving

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

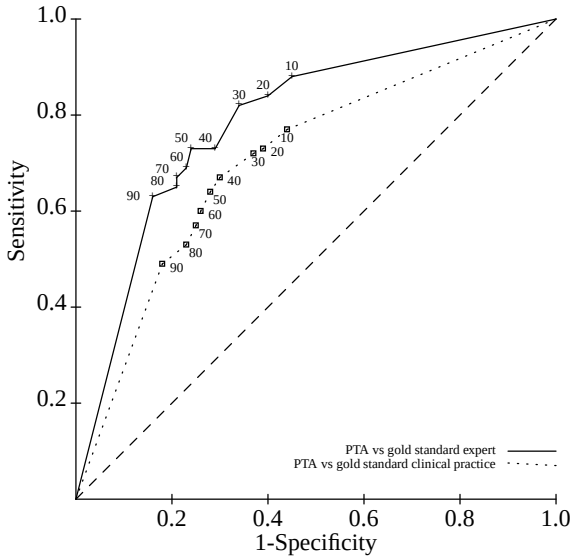

Figure 3: ROC curves based on patient data. The points are labelled with cut-off points in percentages. The upper curve shows that the cut-off point of 50% correlates closest with expert opinion. In relation to the lower plot the upper plot shows that PTA is closer related to expert opinion than to clinical practice.

3.1 Diagnosis of pneumonia

Diagnosing VAP is a difficult task, because none of the symptoms is caused by the pathogens that cause VAP. Especially in the ICU, it is common for a patient to have fever, as there are many causes for fever in addition to pneumonia. Determining a diagnosis based on available evidence $\mathcal{E}$ is usually defined as: $d = \arg\max_{d \in D} \Pr(d \mid \mathcal{E})$, where D here stands for the 'pneumonia' variable, and $\mathcal{E}$ for evidence, such as presence of leucocytosis and body temperature, and duration of hospitalisation and mechanical ventilation. Receiver operating characteristics (ROC) analysis is another frequently used method. It is employed to determine a probability cut-off point, which is then used to establish a diagnosis for future cases [10]. ROC analysis, however, requires a gold standard diagnosis, which often is not available in medicine. This is actually a problem with the diagnosis of VAP, as its pathological diagnosis is very unreliable. The results of an ROC analysis of the model with an infectious disease specialist and the ICU clinicians as gold standards are shown in Figure 3.

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

3.2 Prediction and treatment selection

For the purpose of prediction of likely causative organisms, as well as for the selection of optimal antibiotic therapy, the temporal knowledge incorporated into the Bayesian-network model of VAP is of major importance. Figure 4 clearly indicates that both likelihood of colonisation and pneumonia by particular pathogens vary in time. As in particular the 'colonisation' variables together with selected antibiotics determine choice of treatment by predicting coverage, time cannot be ignored [2].

Treatment selection is nothing else then selecting the antibiotic combination that yields an optimal outcome. In the case of treatment of VAP this can be defined as maximal coverage with minimal side effects. This implies that the Bayesian network needs to be extended with decision theory, i.e. a *utility*

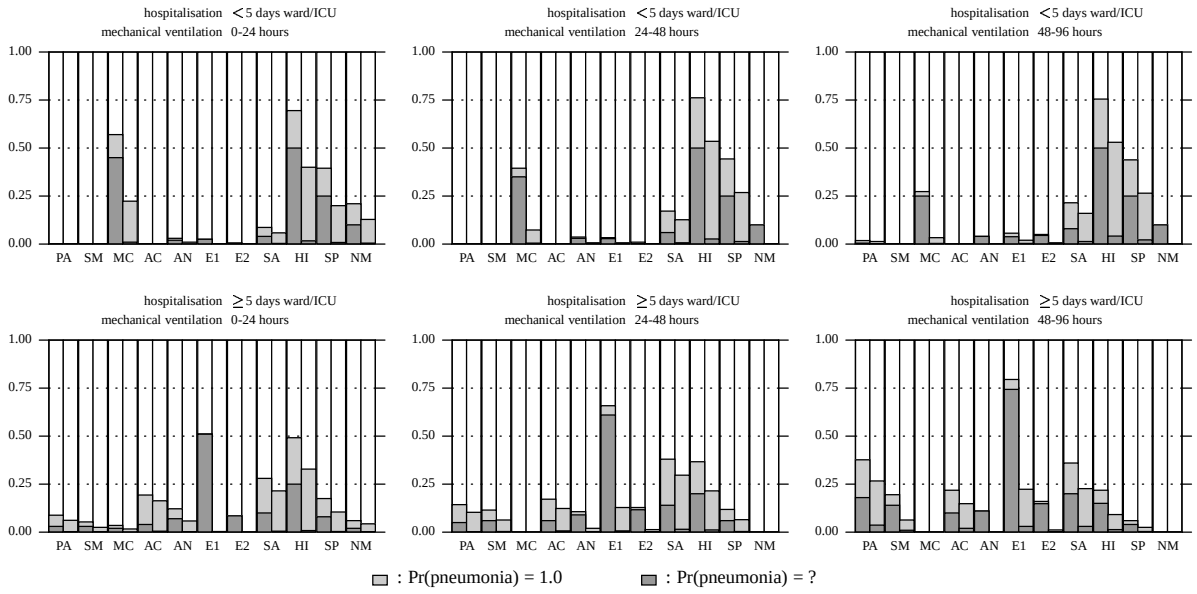

Figure 4: Obtained predictions after entering information concerning duration of hospitalisation and mechanical ventilation. Names of pathogens have been abbreviated. For each pathogen, the probability of colonisation and pneumonia are depicted, in that order.

function u : COVERAGE $\times$ SIDE-EFFECTS $\rightarrow \mathbb{R}$, has to defined and treatment variables become *decision variables*. The resulting formalism is known under various names, among others *decision networks* and *influence diagrams* [8, 9]. The optimal treatment is the one with maximal expected utility.

Influence diagrams can be converted to Bayesian networks, among others by mapping the (bounded) image of the utility function u to the interval $[0, 1]$, and Bayesian-network inference algorithms can be used to determine (the sequence of) optimal decisions. In the VAP model, this mapping is very straightforward, as there is only one decision to make (antibiotic therapy). The actual mapping is derived in Ref. [5].

4 Evaluation

Evaluation of a Bayesian network and decision-theoretic system is not much different from evaluation of any decision-support system. Measures such as true positive and true negative rates can be determines, as accuracy and predictive power for a test dataset with patient data. As Bayesian networks are computational formalisms, it is also possible to carry out cross validation, by leaving out each time a fixed number of cases from a dataset, using the remaining cases to learn the underlying probability distribution, and the selected cases for testing. In addition, the resulting a posteriori probability distributions can be examined using scoring rules, which gives insight into their quality (cf. [6]).

Evaluation of therapy advised by a model usually done by comparing this advice with that of expert clinicians, possibly having treatment prescriptions assessed by an expert panel.

5 Conclusion

In the space available, we have attempted to convey a feeling of the process of developing and using Bayesian-network and decision-theoretic systems which intend to assist medical doctors in diagnosing disease, predicting likely outcome and selecting appropriate treatment. There is now extensive literature available on the technical issues involved, but the amount of literature reporting experience in building actual systems is much more limited. The main reason for this is that it is still a major undertaking to

develop systems for problems of the complexity one usually sees in medicine.

Finally, the main advantages of the BN formalism from a (bio)medical point of view are its flexibility and strong links with how biomedical people think about problems. These properties seem to render BNs one of the most suitable present foundations for medical decision making under uncertainty.

References

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

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

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

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

[5] 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.

[6] P.J.F. Lucas. Expert Knowledge and its Role in Learning Bayesian Networks in Medicine: an Appraisal, in: *AIME 2001*. Lecture Notes in Artificial Intelligence 2101. Springer-Verlag, Berlin, 2001, pp. 156–166.

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

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

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

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