

A Model-based Approach to Improved Prescription of Antibiotics

P.J.F. Lucas¹, C.A.M. Schurink², I.M. Hoepelman^{2,3}, M.J.M. Bonten^{2,3}

¹Institute for Computing and Information Sciences, University of Nijmegen
Nijmegen, The Netherlands, E-mail: peterl@cs.kun.nl

²Department of Medicine, Division of Acute Medicine and Infectious Diseases, and

³Division of Medical Microbiology, Eijkman-Winkler Institute
University Medical Centre, Utrecht, The Netherlands, E-mail: K.Schurink@azu.nl

Abstract

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

1 Introduction

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

In our project *Bayesian networks* (BNs) have been chosen as

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

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

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

2 Modelling

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

¹C2000 is sold by the Eclipsys Corporation, <http://www.eclipsys.com>



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

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

2.1 Ventilator-associated pneumonia

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

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

2.2 The Bayesian-network model

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

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

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

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

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




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

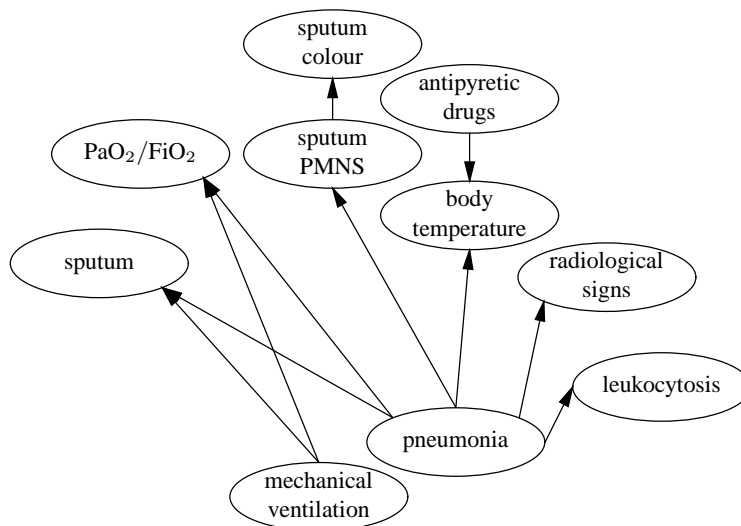


Figure 3: Probabilistic model of signs and symptoms of pneumonia.

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

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

3 Medical problem solving

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

3.1 Diagnosis of pneumonia

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

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

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

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




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

3.2 Prediction and treatment selection

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

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

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

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

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




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




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

4 Implementation and evaluation

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

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

Currently, the system is undergoing a clinical trial. The 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 [4].

5 Conclusion

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

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

References

- [1] J.G. Bartlett. *Management of Respiratory Tract Infections*. Williams & Wilkins, Baltimore, 1997.
- [2] N.C. de Bruijn, P.J.F. Lucas, K. Schurink, M. Bonten, A. Hoepelman. Temporal probabilistic knowledge for medical decision making, in: *AIME 2001*. Lecture Notes in Artificial Intelligence 2101, Springer-Verlag, Berlin, 2001, pp. 231–234.
- [3] R.S. Evans, S.L. Pestotnik, D.C. Classen, T.P. Clemmer, L.K. Weaver, J.F. Orme, J.F. Lloyd, J.P. Burke. A computer-assisted management program for antibiotics and other anti-infective agents. *N Engl J Med* 338(4) (1998) 232–238.
- [4] C.P. Friedman and J.C. Wyatt. *Evaluation Methods in Medical Informatics*. Springer-Verlag, New York, 1997.
- [5] L. Leibovici, V. Geitelman, Y. Yehezkelli, O. Poznanski, G. Milo, M. Paul, P. Ein-Dor. Improving empirical antibiotic treatment: prospective nonintervention testing of a decision support system. *J Intern Med* 242 (1997) 395–400.
- [6] P.J.F. Lucas, H. Boot and B.G. Taal. Decision-theoretic network approach to treatment management and prognosis. *Knowledge-based Systems* 11 (1998) 321–330.
- [7] P.J.F. Lucas, H. Boot and B.G. Taal. Computer-based decision-support in the management of primary gastric non-Hodgkin lymphoma, *Methods of Information in Medicine* 37 (1998) 206–219.
- [8] P.J.F. Lucas, N.C. de Bruijn, K. Schurink, I.M. Hoepelman, A Probabilistic and decision-theoretic approach to the management of infectious disease at the ICU, *Artificial Intelligence in Medicine* 19(3) (2000) 251–279.
- [9] W.J. McIsaac and V. Goel. Effect of an explicit decision-support tool on decisions to prescribe antibiotics for sore throat. *Med Decis Making* 18(2) (1998) 220–228.
- [10] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman, San Mateo, California, 1988.
- [11] S. Renooij. *Qualitative Approaches to Quantifying Probabilistic Networks*. PhD Thesis, Utrecht University, 2001.
- [12] S.D. Shachter. Evaluating influence diagrams. *Operation Research* 34(6) (1986) 871–882.
- [13] J.Q. Smith. *Decision Analysis*. Chapman and Hall, London, 1988.
- [14] H.C. Sox, M.A. Blatt, M.C. Higgins, K.I. Marton. *Medical Decision Making*. Butterworth, Boston, 1988.
- [15] H. Warner Jr, S.R. Blue, D. Sorenson, L. Reimer, L. Li, M. Nelson, M. Barton, H. Warner. New computer-based tools for empiric antibiotic decision support. *JAMIA* (1997) 238–242.