A Probabilistic and Decision-theoretic Approach to the Management of Infectious Disease at the ICU

Peter Lucas
Department of Computing Science, University of Aberdeen
Aberdeen, AB24 3UE, Scotland, UK
E-mail: plucas@csd.abdn.ac.uk

Nicolette de Bruijn
Department of Computer Science, Utrecht University
Padualaan 14, 3584 CH Utrecht, The Netherlands
E-mail: nbruijn@cs.uu.nl

Karin Schurink and Andy Hoepelman
Department of Internal Medicine
University Medical Centre Utrecht
Heidelberglaan 100, 3584 CX Utrecht, The Netherlands
E-mail: {K.Schurink,I.M.Hoepelman}@digd.azu.nl

Abstract

The medical community is presently in a state of transition from a situation dominated by the paper medical record to a future situation where all patient data will be available online by an electronic clinical information system. In data-intensive clinical environments, such as intensive care units (ICUs), clinical patient data are already fully managed by such systems in a number of hospitals. However, providing facilities for storing and retrieving patient data to clinicians is not enough; clinical information systems should also offer facilities to assist clinicians in dealing with hard clinical problems. Extending an information system's capabilities by integrating it with a decision-support system may be a solution. In this paper, we describe the development of a probabilistic and decision-theoretic system that aims to assist clinicians in diagnosing and treating patients with pneumonia in the intensive-care unit. Its underlying probabilistic-network model includes temporal knowledge to diagnose pneumonia on the basis of the likelihood of laryngotracheobronchial-tree colonisation by pathogens, and symptoms and signs actually present in the patient. Optimal antimicrobial therapy is selected by balancing the expected efficacy of treatment, which is related to the likelihood of particular pathogens causing the infection, against the spectrum of antimicrobial treatment. The models were built on the basis of expert knowledge. The patient data that were available were of limited value in the initial construction of the models because of problems of incompleteness. In particular, detailed temporal information was missing. By means of a number of different technique, among others from the theory of linear programming, these data have been used to check the probabilistic information elicited from infectious-disease experts. The results of an evaluation of slightly different models using retrospective patient data are discussed as well.

Keywords: Medical decision support; Probabilistic networks; Bayesian networks; Decision theory; Temporal probabilistic models; Infectious diseases; Intensive care.

1 Introduction

The medical community is presently in a state of transition from a situation dominated by the paper medical record to a future situation where all patient data will be available on-line by means of clinical information systems, also called computer-based patient record (CPR) systems [5]. In many intensive-care units (ICUs) clinical patient data are already fully managed by a clinical information system. Such information systems are potentially interesting sources for clinical research, because they contain large quantities of very detailed patient data. Moreover, these systems offer a natural environment for facilities of decision support that assist clinicians in handling hard clinical problems. The reasons for this are twofold:

(1) Clinicians use these information systems on a daily basis for their patient management, and thus already see information systems as important clinical tools; this view promotes the acceptance of integrated decision-support tools.

(2) Most of the data needed to drive a decision-support system are already available on-line, which makes it possible to employ decision-support systems even under heavy timing constraints.

In the case of the ICU, there are huge quantities of temporal data, like blood-pressure data, available through continuous patient monitoring. Unfortunately, the situation for clinical data concerning the symptoms and signs in the patient and follow-up, is less favourable. These data must be entered by hand, and are usually stored in almost inaccessible free-text form. Yet, these clinical data are normally of crucial importance in the process of medical decision making. Hence, even in hospitals where clinical information systems have replaced the previous paper records, the question remains to what extent such systems may act as a basis for medical decision support.

The present study was undertaken to investigate the potential of a commercial clinical information system (C2000, sold by the Eclipsys corporation) to act as a foundation for medical decision support at the ICU. Although this system had already been in use at three of the four ICUs of the University Medical Centre in Utrecht (UMCU) since 1993, no earlier investigation of this kind had been performed. As a clinical problem domain, the diagnosis and treatment of pneumonia in mechanically-ventilated patients was chosen. It is a problem of major clinical importance to ICUs. A 1994 study of antibiotics usage at Dutch ICUs revealed that 49% of the antibiotics were prescribed for respiratory-tract infections. The problem may, however, also be viewed as an instance of a much wider and even more significant clinical problem: the clinical management of infectious disease in hospitals.

With the advent of modern antimicrobial agents since World War II, it was believed by many that infectious diseases would soon be medical history. However, as new antimicrobial drugs were developed, so did the ability of microorganisms to elude their mode of action. The last two decades, the early feelings of optimism with regard to infectious-disease control have slowly been changing into a feeling of concern [32]. Many scientists now believe that the control of treatment of infectious diseases 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 [12, 13, 24, 34, 44].

The goals of our research project are the design, clinical implementation and evaluation of a decision-support system that is capable of assisting ICU doctors in dealing with patients who are mechanically ventilated and display symptoms and signs possibly related to the development of pneumonia. In this paper, we describe the design of a decision-theoretic model, i.e. a model based on a combination of the theory of (causal) probabilistic networks [23, 15, 31], also known as Bayesian (belief) networks [37], and decision theory [41], that is aimed at supporting clinicians in prescribing antibiotic therapy to mechanically-ventilated patients with pneumonia at the ICU. It is part of a decision-theoretic expert system [25, 27] called *PTA* (Pneumonia Therapy Advisor).

Note that this problem domain is closely related to that of the well-known MYCIN system, an expert system developed in the late 1970s at Stanford University that offered advice on the diagnosis and treatment of sepsis and meningitis [9, 40, 46]. However, the techniques used in PTA are quite different from those that were used in MYCIN: instead of relying on expert classification rules, we employ a decision-theoretic approach which allows for the integration of probabilistic information from various sources, and for improvement of the model based on gathered experience as reflected in clinical databases. Given the complexity of the problem of infectious-disease management, we believe this to be the right approach.

This paper is organised as follows. In the next section, we focus on the diagnosis and treatment of pneumonia, and on the difficulties linked with these issues. Section 3 pays attention to the design of a temporal model of pneumonia. Its implementation in terms of the theory of probabilistic networks is discussed in Section 4, whereas the decision-theoretic extensions to the model are presented in Section 5. Finally, a preliminary evaluation of the models is presented in Section 6. The paper concludes with an analysis of the achievements of this research, a comparison to related research, and future plans.

2 The origin and management of pneumonia at the ICU

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. In contrast, the reported rates of pneumonia developed in the hospital in general, known as *hospital-acquired pneumonia* (HAP) or *nosocomial pneumonia*, are 0.5%–1.0% of all patients admitted [6], which are thus significantly lower than the rates for ICUs. Nevertheless, the associated mortality rates of HAP are high: approximately 30% [21].

In healthy persons, the parts of the respiratory tract distal to the oropharynx are normally sterile due to the activity of a number of defense mechanisms. Two important defense mechanisms are [11]:

(1) *mucociliary clearance*: the beating motion of cilia of the laryngotracheobronchial epithelium moves a film of mucus, which incorporates particles and bacteria that have been inspired, towards the oropharynx thus clearing the epithelium from potential sources of infection;

(2) *alveolar clearance*: bacteria and particles are taken up by alveolar macrophages, which are white blood cells with phagocytic capability, and subsequently digested by means of intracellular enzymes.

Reduced activity of the defense mechanisms results from a number of factors. For example, mechanical ventilation blocks the mucociliary clearance due to the associated tracheal intubation. Since many patients admitted to an ICU need respiratory support by a mechanical ventilator, suppression of this defense mechanism occurs quite often. The capacity of the defense mechanisms may also be surpassed due to reduced consciousness. Many patients in the ICU have a reduction in consciousness. This leads to suppression of the cough reflex, which in turn promotes the aspiration of stomach content with associated pathogens. These factors increase even further the likelihood of pulmonary infection in ICU patients. The integrity of the immune system is another important factor in preventing the occurrence of pneumonia; in patients in which it is suppressed, either congenitally, e.g. an in-born diminished production of immunoglobulin, or acquired, e.g. due to a disease like AIDS, pneumonia is one of the most common causes of death. Among mechanically-ventilated patients, the reported rates of pneumonia are very high: 18%–60%, with an average of about 20% [6]. Pneumonia in mechanically-ventilated patients is called *ventilator-associated pneumonia* (VAP). The reported mortality rates of VAP are also very high: 50%–90% [6].

The treatment of VAP in patients is seen as a significant problem by ICU doctors. Firstly, many of these patients are severely ill. Secondly, 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.

Choosing an appropriate therapy for pneumonia not only involves the issue of susceptibility of pathogens to antibiotic agents, and antimicrobial spectrum, but also of possible side effects of prescribed drugs. In the case of antibiotic therapy possible side effects are: renal failure, diminished hearing, epileptic seizures and allergic reactions varying from skin rash to anaphylactic shock.

What adds very much to the difficulty of the problem is that in many cases it is not even clear whether a patient has ventilator-associated pneumonia or not; the accurate diagnosis of ventilator-associated pneumonia is currently seen as an important clinical challenge. The diagnosis is difficult because of lack of a simple, cheap yet accurate diagnostic test; the disease is therefore diagnosed by taking a number of different clinical features into account [36]. Moreover, patients must usually be treated before the results of sputum cultures become available, which takes at least 48 hours. Sputum cultures yield highly valuable information about the identity and antibiotic susceptibility of pathogens. Hence, the initial therapy, called *empirical therapy*, must be started without having absolute certainty that the patient is affected by ventilator-associated pneumonia and without actually knowing the identity of the causative pathogens.

All factors mentioned above may interact in various manners, in some cases even as a function of time. Clearly, the decisions about appropriate antibiotic therapy must be made on the basis of a lot of uncertain medical knowledge. Offering appropriate decision support to the clinician may therefore be beneficial to the patient; the complexity of the problem is so large that it is very unlikely that clinicians will be capable of delivering optimal treatment to all patients.

3 Temporal modelling of the process of pneumonia

The medical knowledge of pneumonia discussed above was used as a starting point for the design of a number of informal temporal models of ventilator-associated pneumonia. The design was carried out in close collaboration with three infectious-disease experts, as well as

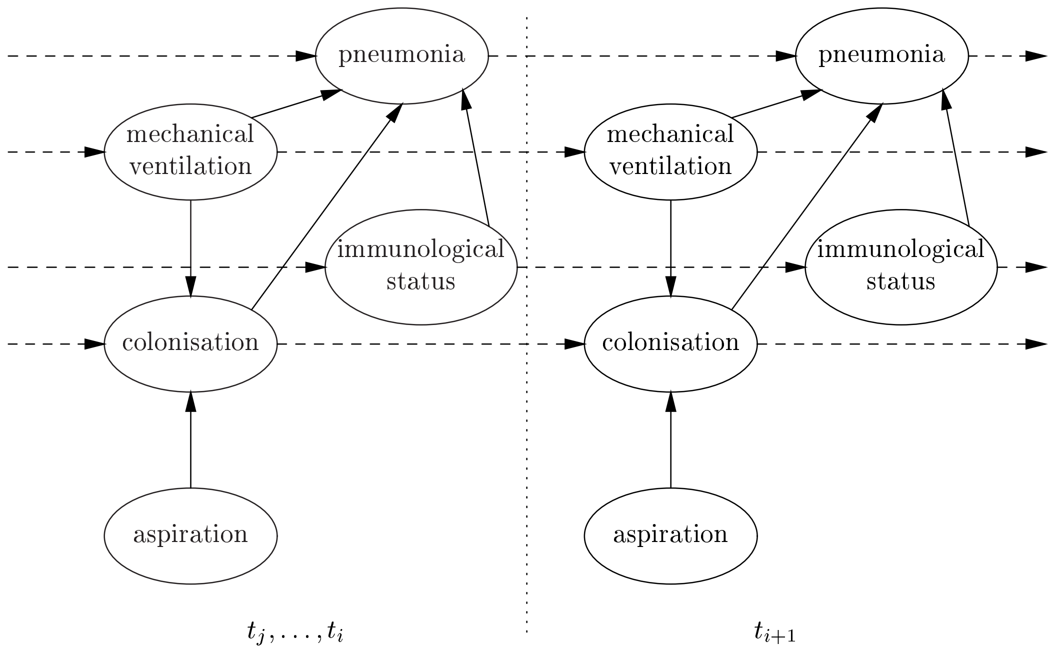

Figure 1: Colonisation as a dynamic process.

with experts from the ICU, from the University Medical Centre Utrecht. We used directed graphs as the main modelling tool in the design process.

As discussed above, the capacity of the defense mechanisms of a patient to pathogens may be diminished, but may also be exceeded. In either case, the laryngotracheobronchial tree and lung parenchyma may become invaded by pathogens, which under normal conditions would be eliminated swiftly. This process is known as *colonisation*. Knowing at any particular point in time that a patient is being mechanically ventilated or has aspirated stomach content, influences our knowledge of whether or not the patient will become colonised by specific bacteria. Hence, there exists a causal relationship between 'mechanical ventilation' and 'aspiration', on the one hand, and 'colonisation', on the other hand. In Figure 1 this causal knowledge has been represented by means of a directed graph with identically named vertices and associated solid arcs, representing causal relationships. Colonisation in turn may give rise to pneumonia. This is expressed in the graph model by a solid arc running from 'colonisation' to 'pneumonia'. Finally, tracheal intubation not only affects the colonisation process, but also directly influences the likelihood that a patient will develop pneumonia. As may be expected, the likelihood that colonisation will give rise to pneumonia is also influenced by the immunological status of the patient. These two independent causal influences have been represented by two solid arcs running from 'mechanical ventilation' and 'immunological status', respectively, to 'pneumonia'.

It is well-known that the identity of the pathogens with which a patient's laryngotracheobronchial tree gets colonised depends on the environment to which a patient is exposed. The pathogens most likely present in a hospital environment are different from those most frequently encountered outside the hospital. For example, *Streptococcus pneumoniae* is the bacterium most likely causing pneumonia outside the hospital, whereas *Pseudomonas aeruginosa* is a likely nosocomial cause of pneumonia. When a patient is already colonised by particular microorganisms when entering the hospital, the pattern of colonisation may change over time.

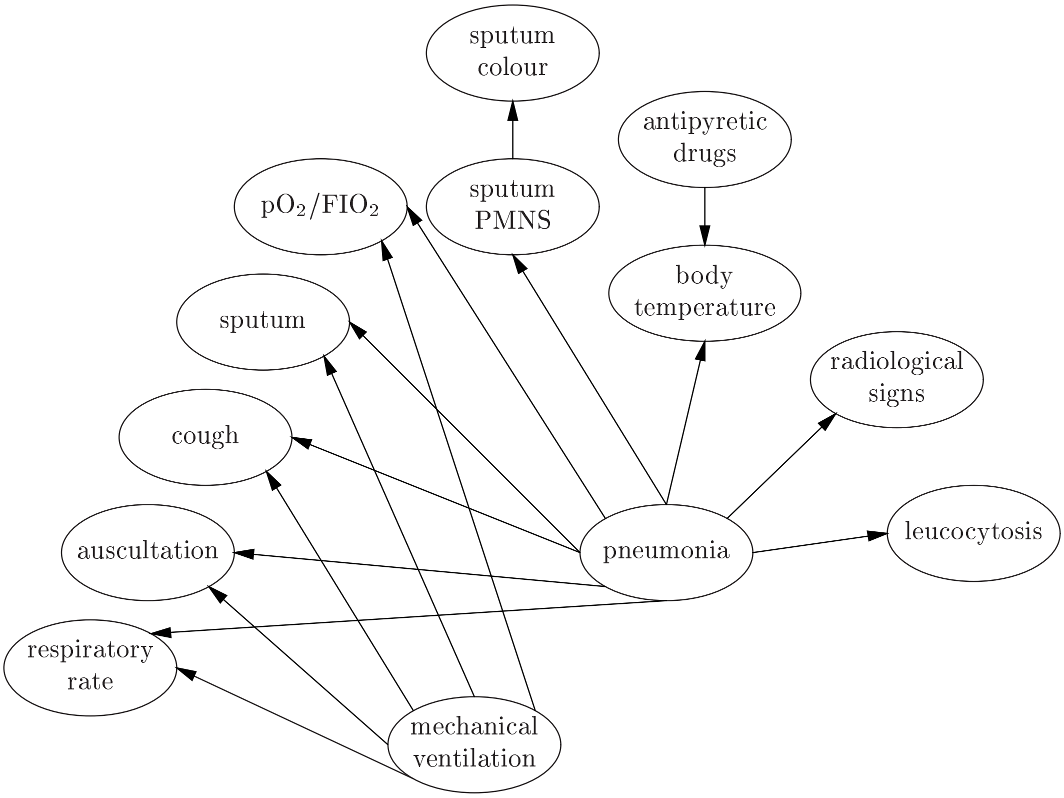

Figure 2: Symptoms and signs of pneumonia.

Hence, colonisation is a dynamic process that evolves over time: only after a certain amount of time, particular microorganisms may be detected in cultures of laryngotracheobronchial secretions.

The temporal nature of colonisation has been indicated in Figure 1 by means of dashed arcs, which have the meaning of temporal causal relationships. Although the process is really a continuous one, here it is assumed that the process is discrete, since data about laryngotracheobronchial-tree colonisation will only be available at discrete time points from examined sputum cultures. The vertices in the graph may now be viewed as *state variables*. Moving to a similar state at the next point in time is said to be moving to the next *stage*. Note that it is assumed that our knowledge of every state, with the exception of 'aspiration', at time point t_{i+1} is influenced by knowledge of similar states at time points $t_j, \ldots, t_i$, i.e. by $i - j + 1$ previous stages. For example, when we know that a patient has been mechanically ventilated for the last two days ($j = i - 1$), the likelihood that the patient will remain mechanically ventilated is different from when we had assumed that the patient was, for example, mechanically ventilated over the last three days ($j = i - 2$). When $j = i$, it is assumed that only the previous stage should be taken into account when predicting the next stage. The actual value of j, however, is not easy to establish, and it seems reasonable to assume that the immediate previous stage is the one most important in predicting the next stage. Henceforth, it is assumed that $j = i$. Furthermore, it is assumed that knowing that a patient has aspirated stomach content at a particular stage does not influence the likelihood that aspiration occurs in the next stage. This is only correct when there is no systematic cause underlying aspiration.

Associated with pneumonia are particular symptoms and signs that can be employed to diagnose the presence of pneumonia in a patient. The nine most significant symptoms and signs are indicated in Figure 2 as effects of a common cause 'pneumonia'. Some of these are suppressed when a patient is mechanically ventilated. This explains why some effect vertices have a second incoming arc from the vertex 'mechanical ventilation'. Furthermore, the body temperature, taken as an indication of presence or absence of fever in a patient, is influenced by taking antipyretic (anti-fever) drugs. The variable pO_2/FIO_2 represents the ratio between arterial oxygen pressure and the fractional inspired oxygen level. The variable 'sputum PMNS' represents the number of polymorphonuclear leucocytes found in a sputum sample examined under the microscope.

To predict whether or not a particular antimicrobial therapy may improve a patient's condition, it is necessary to know the susceptibility of the suspected pathogens to particular antibiotics. The basic temporal model has therefore been extended as indicated in Figure 3. Determining antibiotic coverage of the pathogens possibly causing pneumonia is the main goal of the system, which explains the inclusion of a vertex 'coverage'. The vertex 'side effects' indicates that some antibiotics may give rise to side effects at a particular point in time; these are also predictive for later stages. Note that it is assumed that both 'susceptibility' and 'coverage' do not influence future susceptibility and coverage: susceptibility and coverage are fully determined by the nature of the microorganisms with which a patient is colonised.

4 A probabilistic model of pneumonia

The temporal model discussed above offers a meticulous representation of the structure of the processes underlying the development and treatment of pneumonia. However, we have not yet dealt with the many uncertainties involved in these processes, which, as was mentioned above, is one of the main issues of the problem. It will be the subject of the present section.

4.1 The probabilistic-network formalism

A natural representation of the uncertainties involved in treating patients with pneumonia is offered by the probabilistic-network formalism [31, 37]. This formalism allows for the representation of causal relationships, either explicitly indicated as being temporal in nature or not. Such relationships were used in the design of the dynamic model discussed above.

Formally, a *probabilistic network* is an acyclic directed graph $G = (V(G), A(G))$, with a set of vertices $V(G) = \{V_1, \ldots, V_n\}$, where each vertex $V_i \in V(G)$ represents a discrete stochastic variable, and a set of arcs $A(G) \subseteq V(G) \times V(G)$ reflecting all known stochastic dependencies in the domain concerned. Arcs are often informally seen as to mirror causal or correlational influences among variables. Absence of arcs between vertices reflects known (conditional) independencies among stochastic variables. Stochastic variables will be denoted by upper-case letter, e.g. X; values of variables will be denoted by lower-case letters, e.g. x. In the case of binary variables, the value $X = \mathit{yes}$ is sometimes simply denoted by x; the value $X = \mathit{no}$ is also denoted by $\neg x$. By $\sum_X \varphi(X)$ is indicated a summation over the values of the variable X of the function φ. Values or *states* of specific variables, such as COL (colonisation), are also indicated using this notation.

On the set of variables $\{V_1, \ldots, V_n\}$ is defined a joint probability distribution Pr that can

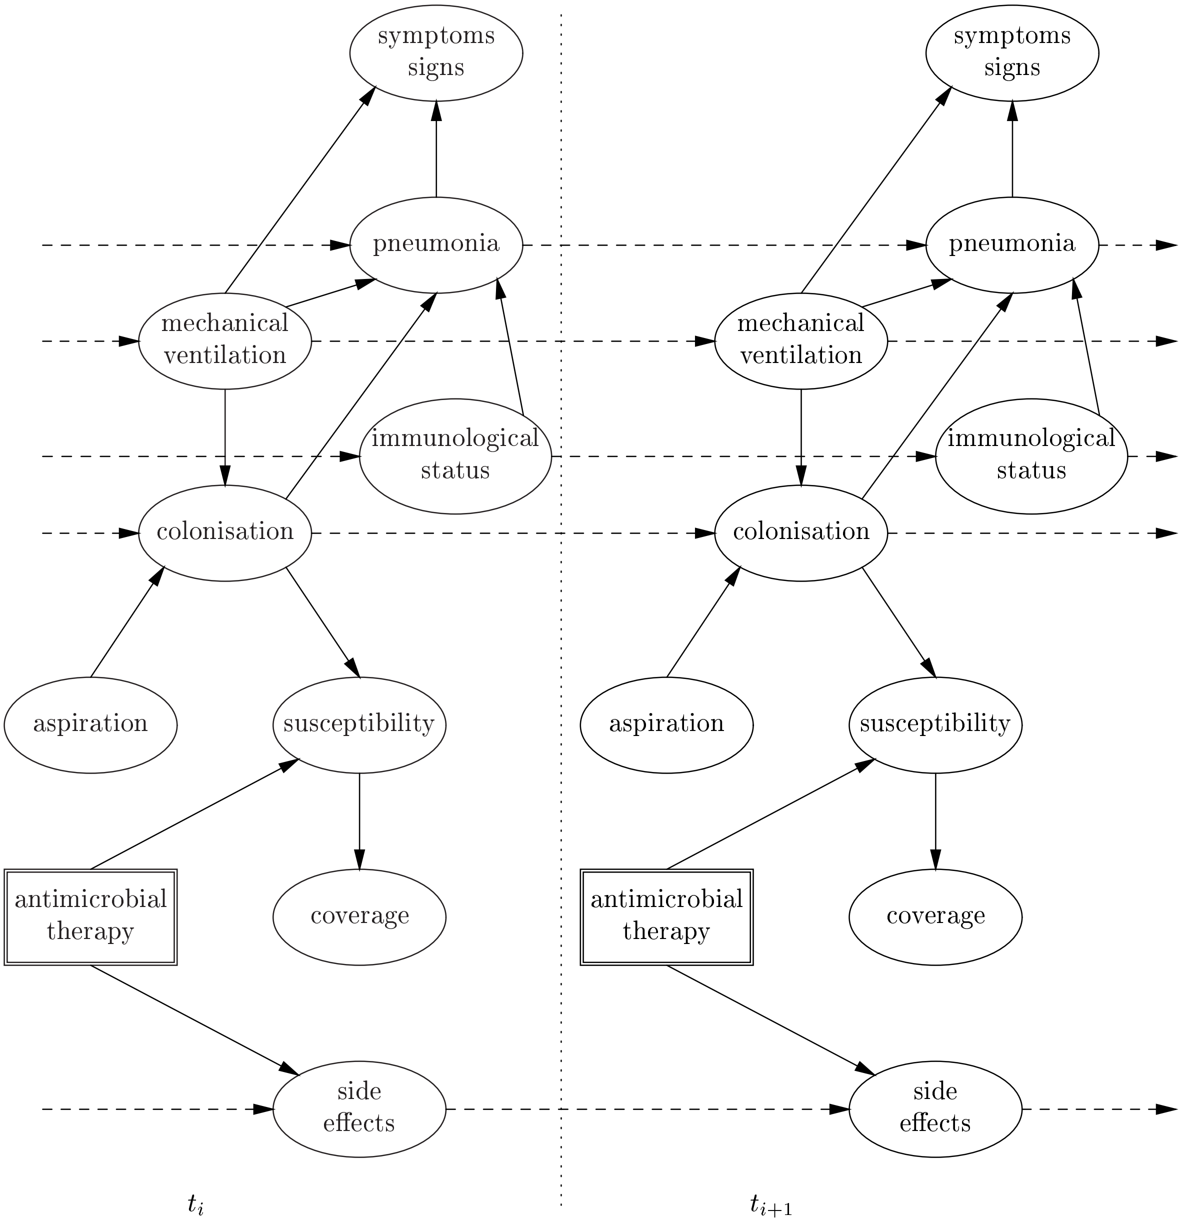

Figure 3: Complete temporal model.

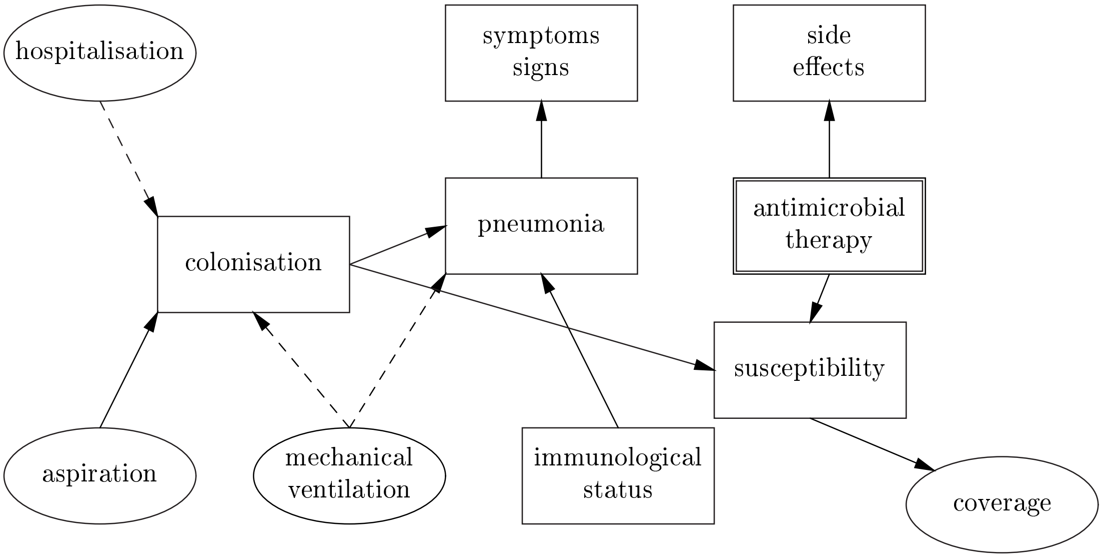

Figure 4: Global structure of the probabilistic VAP model.

be factorised according to the topology of the graph as follows:

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

where $\pi(V_i)$ represents the set of variables associated with the parent vertices of V_i. This means that the joint probability distribution $\Pr(V_1, \ldots, V_n)$ can be defined in terms of 'local' probability tables $\Pr(V_i \mid \pi(V_i))$ by assuming the variable V_i to be conditionally independent of all predecessors of the associated vertex V_i given the parents $\pi(V_i)$.

The probabilistic-network formalism also offers facilities for entering and processing evidence E, using probabilistic-inference algorithms [20, 23, 37]. A probabilistic-inference algorithm updates a given probability distribution Pr yielding a new probability distribution $\Pr^*$ that takes the given evidence into account. The following equality holds between these two probability distributions:

$$\Pr^*(V_i) = \Pr(V_i \mid E) \tag{1}$$

i.e. the marginal probability distribution of any variable V_i in the updated probability distribution $\Pr^*$ reflects the original probability distribution Pr for the variable V_i when conditioned on the given evidence E.

4.2 The VAP model

Above we have sketched the development of pneumonia in a patient as a dynamic process. The adopted view was that colonisation and related infectious-disease processes go through various stages, with every subsequent stage fully determined by the previous stage. When the vertices in the dynamic model are assumed to stand for stochastic variables, this assumption is known as the *no-memory property* or *Markov assumption* [16, 38]. It holds for example that:

$$\Pr(\textsc{col}_{t_{i+1}} \mid \textsc{col}_{t_0}, \ldots, \textsc{col}_{t_i}) = \Pr(\textsc{col}_{t_{i+1}} \mid \textsc{col}_{t_i})$$

where 'COL' stands for colonisation with some type of microorganism. When we adopt this assumption and take the structure of the temporal model shown in Figure 3, a Markov process results [8, 16, 38].

A major disadvantage of a dynamic model like the one depicted in Figure 3, however, is that it is very demanding not only with regard to the amount of data needed to specify the probability distribution underlying the stochastic process, but also computationally [8]. It is therefore usually necessary to use additional assumptions to further simplify the model. The model that resulted from this effort is shown in Figure 4. Boxes in this model indicate collections of similar stochastic variables, with the therapy variable as a special double-box case; ellipses indicate single stochastic variables. Dashed arcs denote *temporal* probabilistic relationships; solid arcs represent stochastic dependency without special temporal meaning.

As before, central to the model is the temporal process of *colonisation* of the laryngotracheobronchial tree by pathogens. The temporal nature of the process is expressed by the interaction between duration of stay at the ICU (*hospitalisation*) and the duration of mechanical ventilation: both duration of stay at the ICU and duration of mechanical ventilation are correlated to colonisation by pathogens. This part of the model hides the temporal nature of the development of the processes of colonisation and pneumonia in conditioning variables, whereas time was previously handled more explicitly in the model shown in Figure 3.

The relationship between 'aspiration' and 'colonisation' remains the same, as is true for the symptoms, signs or laboratory abnormalities that can be observed, which are once more summarised in a corresponding vertex in the graph. Finally, the susceptibility of pathogens to particular antimicrobial treatment is determined by the choice of medical treatment and the pathogens actually present, causing infection.

As mentioned above, the box-shaped vertices shown in the graph in Figure 4 actually comprise a number of separate, but similar vertices. For example, colonisation by pathogens was modelled as a biological process, in which it was assumed that colonisation by different pathogens occurs independently. This is shown in Figure 5. Note that this representation allows for a pulmonary infection to occur in a patient due to *multiple* organisms. Colonisation by seven of the most frequently occurring pathogens, such as *Pseudomonas aeruginosa* and *Haemophilus influenzae*, is represented in the current model. The number of pathogens included in the model has decreased in time from eleven in an earlier model, to seven in the present model, after a decision was made to focus on empirical treatment only. Similarly, the vertex 'susceptibility' represents the separate effect of antimicrobial treatment on each of these seven pathogens.

To model the probabilistic interaction of the various pathogens on the symptoms and signs of pneumonia, the notion of *causal independence* [17] was used. Consider the probabilistic network shown in Figure 6. Assuming that the probability distribution $\Pr(E \mid I_1, \ldots, I_n)$ that is specified for the variable E expresses some deterministic function $f : I_1 \times \cdots \times I_n \to E$, the probability of the effect $E = e$ given the causes $C_1, \ldots, C_n$ can be computed straight from the axioms of probability theory, as follows. According to Figure 6, the causes C_j are assumed to be mutually independent and the variable E is conditionally independent of any C_j given $I_1, \ldots, I_n$. It follows that:

$$\Pr(e \mid C_1, \ldots, C_n) = \sum_{I_1, \ldots, I_n} \Pr(e \mid I_1, \ldots, I_n) \prod_{k=1}^{n} \Pr(I_k \mid C_k) \qquad (2)$$

Now, the probability distribution on the variable 'pneumonia' models the disjunctive effect of different pathogens, assuming that pneumonia may even be caused by one type of microor-

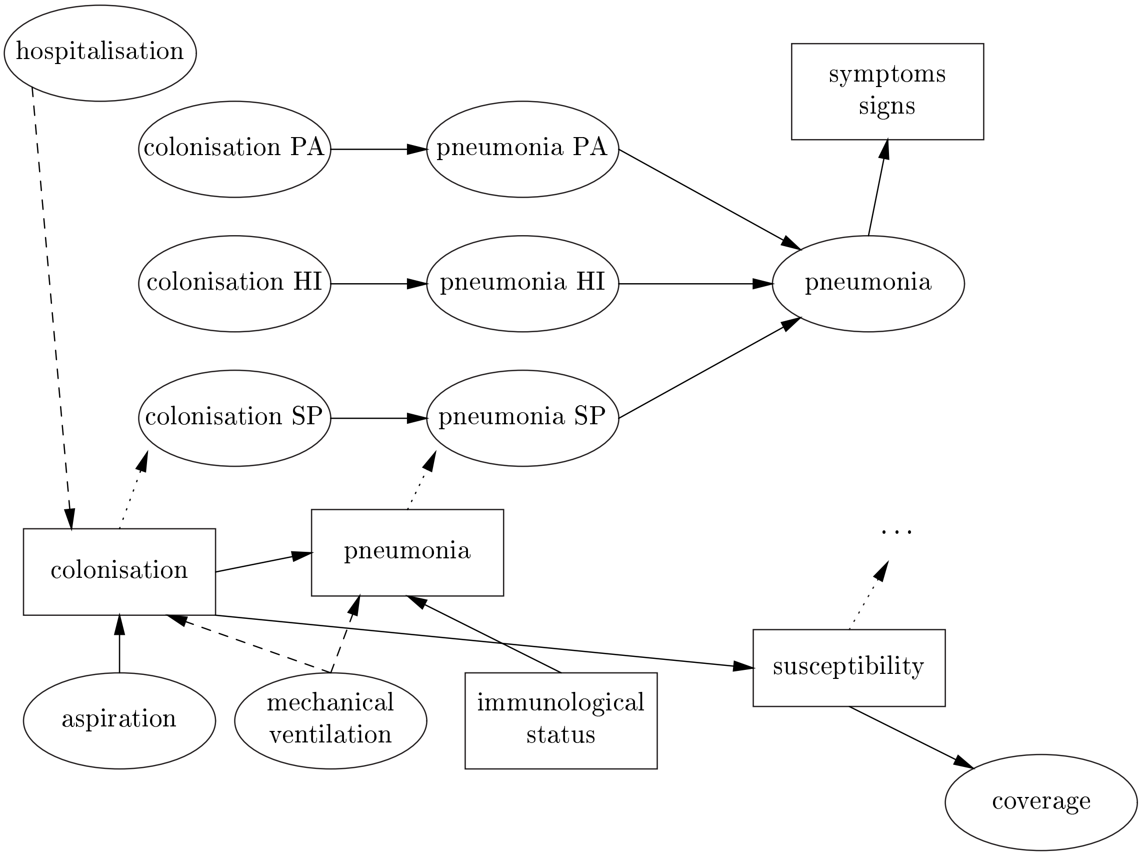

Figure 5: Detailed structure of part of the probabilistic VAP model. Only three of the seven microorganisms included in the model are shown. Dotted arcs point to the actual topology of the probabilistic network. Abbreviations: PA = Pseudomonas aeruginosa, HI = Haemophilus influenzae, SP = Streptococcus pneumoniae.

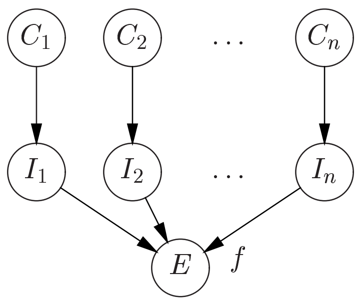

Figure 6: Causal independence model.

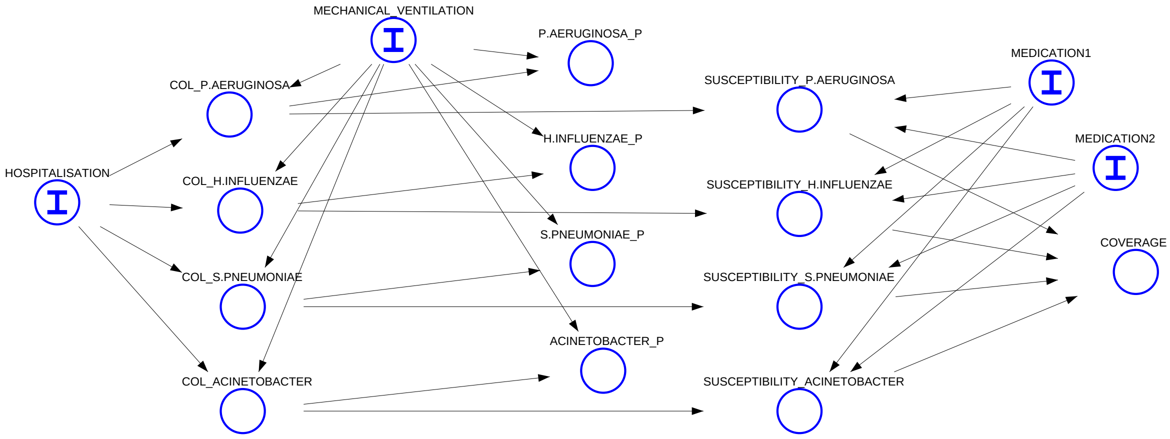

Figure 7: Part of the probabilistic VAP model.

ganism. This principle is modelled by a probability distribution $\Pr(E \mid I_1, \ldots, I_n)$ that is defined as a logical OR:

$$\Pr(e \mid I_1, \ldots, I_n) = \begin{cases} 1 & I_j = i_j \text{ for some } j,\, 1 \leq j \leq n \\ 0 & \text{otherwise} \end{cases}$$

As a consequence, we have

$$\Pr(\text{pneumonia}) = 1 - \prod_{i=1}^{n} \Pr(\neg\text{col-organism}_i)$$

This probabilistic model is known as the *noisy-OR gate* [17, 26, 37]; it is used quite frequently in medicine. A similar type of causal independence was used to model the conjunctive effect of antibiotics on the susceptibility of pathogens. Here, the probability distribution defined on the variable COVERAGE was taken to represent a logical AND; formula (2) now yields the following result:

$$\Pr(\text{coverage}) = \prod_{i=1}^{n} \Pr(\text{susceptibility-organism}_i)$$

which is known as the *noisy-AND gate* [17, 26]. In Figure 7, a small part of the model as implemented using an improved version of the Ideal probabilistic expert-system shell [43], is shown, only involving the vertices *hospitalisation*, *mechanical ventilation*, *colonisation* (e.g. `COL_P.AERUGINOSA`, i.e. colonisation by *Pseudomonas aeruginosa*), and *pneumonia* (e.g. `P.AERUGINOSA_Pneumonia`), *susceptibility* of pathogens to antimicrobial drugs, and *medication*, and restricted to four of the seven included pathogens. Vertices with capital `I` are *instantiation vertices*, i.e. variables without parents that must always be supplied with a value. When a value for the corresponding variable is supplied, additional stochastic independence information is yielded. In this way, the speed of probabilistic inference could be dramatically improved.

Antibiotic treatment consists of the selection of one or two different antibiotics – possibly also none – modelled by two identical therapy vertices; each therapy vertex includes 24 different antibiotic agents (including none), yielding $24^2 = 576$ possible combinations, of which

$\binom{23}{2} = 253$ (excluding 'none') are different, yielding a total of $253 + 24 = 277$ (now including 'none' and single drugs) different therapies.

4.3 Probability assessment

Although the three ICUs that acted as setting of the research all use the same shared computer-based patient record system, and stopped using paper patient records six years ago, it appeared very hard to select relevant patient cases from the collected databases. The main reason was that ventilator-associated pneumonia is always a concomitant disease. Therefore, clinicians do not find it worthwhile to report the presence of VAP in a patient. We found that only in a very small proportion of cases, much smaller than expected from the reported rates in the literature, patients had been affected by ventilator-associated pneumonia. The infectious-disease specialists also expected much higher rates. As a result of this, all probabilistic information in the model had to be assessed using expert estimates, and, when possible, compared and adjusted to information found in the clinical literature.

There were also a number of databases available from the microbiology laboratory with data about sputum cultures of ICU patients. Because not all variables included in the model were present in these databases, the data could not be used to assess the probability distribution of colonisation by pathogens. We have, however, tried to use the database to *check* the expert estimates for the 'colonisation by pathogen' variables, without any doubt the most crucial variables in the model. Below, we discuss the methodology we have developed to check probabilistic information, using feature-incomplete data. In particular, we shall focus attention on the necessary assumptions to utilise such databases.

Figure 8 shows the percentage of patients for each day at one of the ICUs that had been staying at that particular day less than 5 days. To obtain this curve, it was necessary to make the assumption that all beds at the ICU are always occupied, because sputum cultures are not taken for all patients, and often the number of patients present in the ICU at one particular day seemed unrealistically low. For the 'missing' patients, it was assumed that they were staying less than 5 days, which seems a reasonable assumption.

The reason for utilisation of the database was to check the probability distribution:

$$\Pr(\textsc{col-organism}_i \mid \textsc{hosp}, \textsc{mv})$$

For the purpose of illustration, we restrict our attention to colonisation by *Pseudomonas aeruginosa* (COL-PA, with 'COL-PA = yes' abbreviated to 'col-pa'), hospitalisation (HOSP) either less than 5 days for patients with (lt5days-copd) or without (lt5days-icu) COPD (Chronic Obstructive Pulmonary Disease, like chronic bronchitis), and mechanical ventilation (MV) between 24 and 48 hours. From the axioms of probability theory, it follows that:

$$\begin{aligned}
&\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days}) = \\
&\qquad \Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu} \vee \textsc{hosp} = \text{lt5days-copd}) = \\
&\qquad [\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu}) \Pr(\textsc{hosp} = \text{lt5days-icu}) + \\
&\qquad \Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-copd}) \Pr(\textsc{hosp} = \text{lt5days-copd})]/ \\
&\qquad \Pr(\textsc{hosp} = \text{lt5days})
\end{aligned}$$

Note that $\Pr(\textsc{hosp} = \text{lt5days}) = \Pr(\textsc{hosp} = \text{lt5days-copd}) + \Pr(\textsc{hosp} = \text{lt5days-icu})$, because the values of the variable HOSP are mutually exclusive.

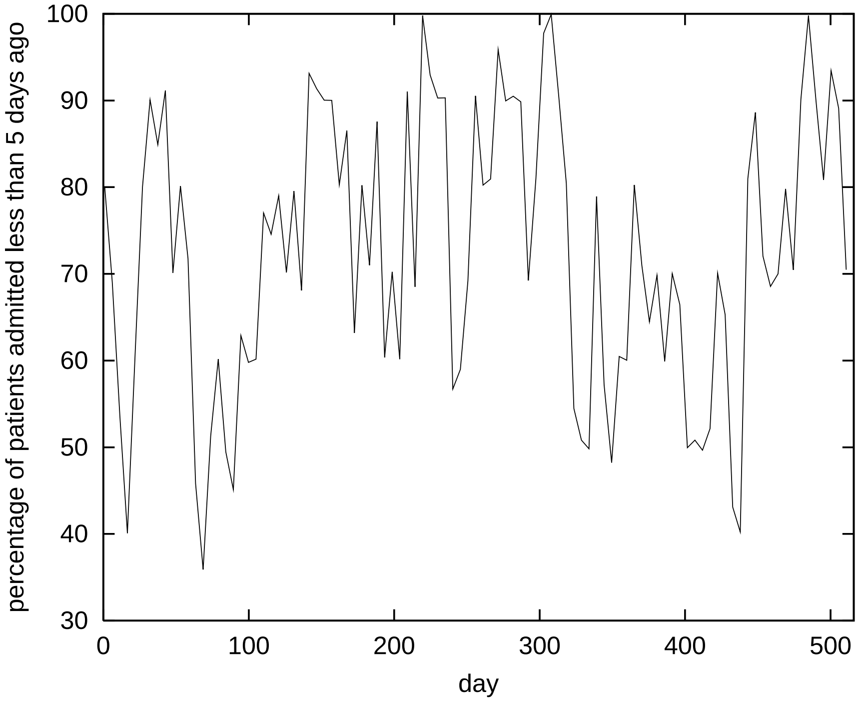

Figure 8: Percentage of patients in one of the ICUs being admitted to the ICU less than 5 days ago from 1 January 1997 to 1 June 1998, based on a sputum-culture database. It was assumed that all beds are always occupied. The curve was obtained by cubic spline interpolation. Note that the percentage of patients that have been staying longer than 5 days is obtained by taking the x-axis image of the graph.

From the database with data from one of the ICUs it was calculated that:

$$\begin{aligned}\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days}) &= 0.1142\\ \Pr(\textsc{hosp} = \text{lt5days}) &= 0.7207\end{aligned}$$

Furthermore, we assumed that the number of patients with COPD at the ICU is $\frac{1}{5}$th of patients without COPD and who were staying less than 5 days at the ICU, which seemed a reasonable estimate:

$$\Pr(\textsc{hosp} = \text{lt5days-copd}) = 0.2 \cdot \Pr(\textsc{hosp} = \text{lt5days-icu})$$

Using linear programming, the following equality was obtained:

$$\begin{aligned}&\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu}) =\\ &\quad -0.2 \cdot \Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-copd}) + 0.13704\end{aligned}$$

with $\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-copd}) \in (0.1142, 0.6852)$.

Next, it was assumed that $\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu}, \textsc{mv} = \text{24-48})$ was a linear function of time t in terms of $\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu})$:

$$\begin{aligned}&\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu}, \textsc{mv} = \tau) \;=\\ &\quad \Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu}) + (1 - \Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu})) \cdot \tfrac{t}{t_{\max}}\end{aligned}$$

with $t \in \{0, 24, 48, 96\}$ and τ = 0-24, 24-48, 48-96, 96-144, > 144, and t the lower bound of τ; for example, when τ = 24-48, then $t = 24$. For $t_{\max}$ we took values ≥ 480 h. Now, when the expert estimate $\Pr(\text{col-pa} \mid \textsc{hosp} = \text{lt5days-icu}, \textsc{mv} = \text{24-48}) = 0.05$ was added to the equations and and inequalities above, the resulting system appeared to be consistent.

We conclude that many assumptions have to be made when checking probability estimates against a database that lacks information about conditioning variables. However, the method discussed above forces one to look more closely at these probability estimates, also in comparison to other probability estimates, thus facilitating probability refinement.

5 Decision making

Until now, we have focussed on the probabilistic model underlying PTA. In this section, a number of extensions to render the model suitable for selecting optimal antibiotic therapy will be paid attention to.

5.1 Representation and problem-solving

The aim of our decision-theoretic expert system PTA is to aid clinicians in dealing with patients with suspected VAP. Insight into the potential efficacy of treatment of such patients can be obtained by entering symptoms and signs of a patient, duration of hospitalisation and mechanical ventilation, previously experienced side effects of antibiotics, and relevant laboratory data into the probabilistic network. Only data concerning vertices that are denoted as instantiation vertices need to be supplied; entering more patient data, however, may provide additional evidence for the presence or absence of pneumonia, but the model is still applicable when most other findings are unknown. Using one of the probabilistic-inference algorithms (cf. [10, 20, 23, 37]), the probability distribution defined on the COVERAGE variable will be updated, yielding a marginal probability summarising the coverage of all possible pathogens $\Pr^*$ reflecting the entered evidence (*Evidence*), i.e.

$$\Pr^*(\textsc{coverage}) = \Pr(\textsc{coverage} \mid \mathit{Evidence})$$

Automatic selection of optimal treatment for a patient requires extending the probabilistic network to a *decision-theoretic model*. A decision-theoretic model not only includes uncertain probabilistic knowledge, but also the preferences among combinations of decisions and states, expressed by means of a utility function, which guide the choice among the various decisions. Often, *influence diagrams* are used for the representation of both probabilistic knowledge, decisions and utility information; an advantage of this formalism is that its probabilistic part can still be viewed as a probabilistic network [39].

We have designed a number of different utility models to enable determining optimal therapy for a patient. A straightforward example of a utility model is a utility function $u : \textsc{coverage} \to \mathbb{R}$, which takes its maximum value when $\textsc{coverage} = \text{yes}$ and its minimum value when $\textsc{coverage} = \text{no}$. Optimal coverage can be computed by varying the therapy choice, and selecting the therapies with maximum expected utility. However, since the number of possible drug combinations was 576, computation of optimal therapy was practically not feasible. As a practical solution to this problem, we have restricted the 277 different therapies to the 33 different therapies considered adequate for most patients; this therapy choice was then represented as a single vertex.

As discussed in the Introduction, prescribing antibiotics is essentially a trade-off between maximising coverage and minimising broadness of spectrum. When experimenting with the model and a database of retrospective patient data (See Section 6.1), it became apparent that the spectrum of the antibiotics suggested by the system was often too broad. This was consistent with our expectations. Two new utility functions of the form u_i : COVERAGE $\times$ SPECTRUM $\rightarrow \mathbb{R}$, $i = 1, 2$, were therefore subsequently designed. For utility function u_1 (*utility model I*) antibiotics or combinations of antibiotics were classified into five different groups from very narrow to very broad according to their spectrum. Utility function u_2 (*utility model II*) was similar to u_1, except that the antibiotics were classified into eight different groups. In both cases, the design of the utility model was done with the help of the direct scaling method [42]. Utility models are often dependent on personal preferences, and this holds especially when designing multi-attribute utility models; no correct utility model exists beforehand. Only evaluation will tell whether or not one model is better than another, as will be discussed below.

As mentioned above, the probabilistic network underlying the decision-theoretic model was structured in such a way, as to cover all pathogens in proportion to their likelihood of occurrence in a patient. The noisy-AND probabilistic model was the main vehicle for this approach. After much experimentation, however, it was suspected that the noisy-AND model had a tendency to broadening the spectrum of the antibiotics advised for prescription. As an alternative, a problem-solving method was devised in which only pathogens with a posterior marginal probability larger than 30% were taken into account; the remaining pathogens were ignored by setting their marginal probability to 0. The threshold of 30% was chosen by the infectious-disease expert after thorough examination of the results of the model for a number of patients. The hypothesis was that by combining utility model I or II with probabilistic threshold problem-solving, the spectrum of the antibiotics advised for prescription would narrow down. The results of an evaluation study in which this hypothesis is tested, are described in Section 6.2.

5.2 Decision-theoretic computation

By means of special decision-theoretic algorithms, such as the algorithms by Shachter [39] and Cooper [10], it is possible to determine the sequence of decisions yielding a maximum expected utility. In the present case, however, it appeared straightforward to use instead a probabilistic-inference algorithm as a basis for the decision-theoretic computations. Consider the influence diagram shown in Figure 9(a), which offers an abstract representation of the decision-theoretic VAP model, i.e. the influence diagram obtained by taking the 'medication' vertex as a *decision vertex*, represented as a box, and by adding a *value vertex*, representing a utility function, which is denoted as a diamond. In the figure, D is a decision variable and U denotes the value vertex with associated utility function. The influence diagram also includes two chance vertices, represented as circles, standing for the stochastic variables E and B. These variables may, in fact, also be taken to stand for probabilistic submodels.

Let $u : D \times B \rightarrow [l, r]$, with $l, r \in \mathbb{R}$, $l < r$, be a surjective utility function defined for the value vertex U. The expected utility of decision $D = d$, $\text{eu}(d)$, is then equal to:

$$\text{eu}(d) = \sum_{B} u(d, B) \cdot \Pr(B \mid e, d) \tag{3}$$

where e is evidence entering for the variable E into the network. A decision with maximum

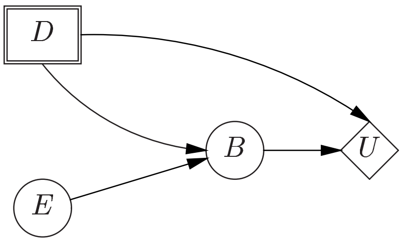

(a) Influence diagram

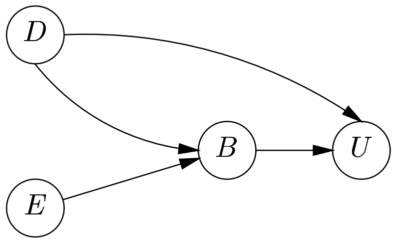

(b) Probabilistic network

Figure 9: Transformation of influence diagram to probabilistic network.

expected utility, $d_{\max}$, yields a maximum value for $\mathrm{eu}(d)$:

$$d_{\max} = \arg\max_{d \in D} \mathrm{eu}(d)$$

Now the same result can be obtained by using a probabilistic network corresponding to the influence diagram, as shown in Figure 9(b), where the probability distribution $\Pr(u \mid B, D)$ is defined as follows:

$$\Pr(u \mid B, D) = u(D, B)/(r - l)$$

The complementary probability $\Pr(\neg u \mid B, D) = 1 - \Pr(u \mid B, D)$ is actually ignored in the decision-making process.

The posterior probability distribution $\Pr(u \mid e, d)$ can now be obtained as follows (it is equal to the marginal probability $\Pr^*(u)$ obtained from a probabilistic-inference algorithm):

$$\Pr(u \mid e, d) = \sum_B \Pr(u, B \mid e, d) \quad (4)$$

$$= \sum_B \Pr(u \mid B, d) \cdot \Pr(B \mid e, d) \quad (5)$$

because U is conditionally independent of E given B. As can be concluded from the equations (3) and (5), it holds that:

$$d_{\max} = \arg\max_{d \in D} \Pr(u \mid e, d)$$

because $\mathrm{eu}(d) = (r - l) \cdot \Pr(u \mid e, d)$, and we can thus use a probabilistic-inference algorithm with the probabilistic network obtained from this straight-forward transformation to compute optimal decisions. We have done so accordingly.

6 Evaluation

The present structure of the probabilistic network has a strong logical foundation, and we therefore believe it to be basically correct. There may be particular arcs missing due to gaps in the medical knowledge concerning pneumonia. However, correlations due to missing medical knowledge are likely to be weak, and hence have little effect on the probability distribution.

In addition to the structure of a probabilistic network, the accuracy of the represented probabilistic and utility information is an issue that requires attention [35, 45]. This has been investigated in a number of different ways, as will be discussed below.

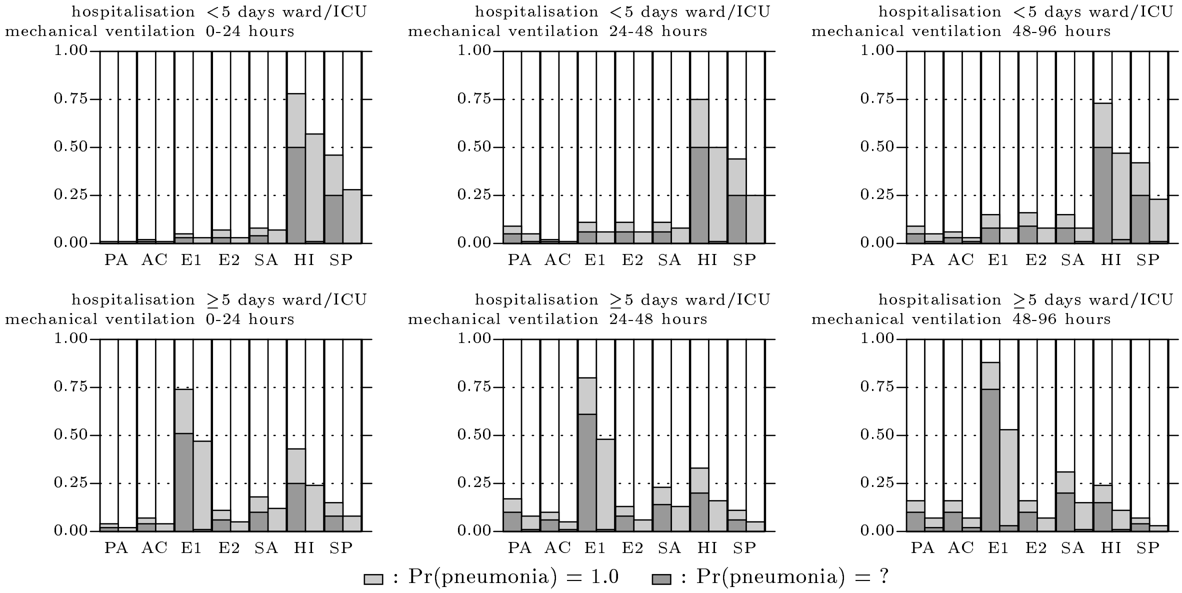

Figure 10: Obtained predictions after entering information concerning duration of hospitalisation and mechanical ventilation. Names of pathogens have been abbreviated. For each pathogen, the probability of colonisation and pneumonia are depicted, in that order.

6.1 Preliminary study

In a preliminary study, the behaviour of the probabilistic model was examined by considering changes in probabilistic patterns as a function of time, and by reviewing the results for a number of individual patients.

First, some of the vertices in the model were instantiated (i.e. a specific value was chosen), and the resulting posterior probability distribution was compared with frequency of occurrence information available from the ICU; see Figure 10. It appears that the patterns of frequency of colonisation and pneumonia for particular pathogens changes with duration of stay at the hospital. For example, the relative frequency of colonisation and pneumonia by *Haemophilus influenzae* (HI in the figure), which is high when entering the hospital, decreases just a bit during the first 96 hours of mechanical ventilation (upper three graphs), to become even lower when the patient has stayed more than five days in hospital. This frequency decreases further for patients who are mechanically ventilated during the next 96 hours (lower graph). In contrast, the relative frequency of colonisation and pneumonia from *Pseudomonas aeruginosa* (PA in the figure) rises after a stay of more than five days in the ICU, which is increased even further due to mechanical ventilation. These observations did agree with expert opinion.

Furthermore, we have examined the results produced by the model for a number of real patients. We discuss here the results for two of these patients in order to convey how the model might be employed clinically, and to point out some of the limitations of the basic utility model that was purely meant to optimise coverage.

Patient A was already mechanically ventilated in the ICU for two days, and had been so since admission to the ICU. The patient produced purulent sputum, and auscultation of the lungs revealed the presence of rales. However, there were no radiological signs of pneumonia,

Predictions/ Advice	**Patient A**	Pr	**Patient B**	Pr
Colonisation (predicted)	H. influenzae S. pneumoniae S. aureus P. aeruginosa	49% 24% 6% 5%	H. influenzae S. pneumoniae S. aureus P. aeruginosa	70% 40% 19% 14%
Sputum cultures	H. influenzae		H. influenzae S. aureus	
Pneumonia	no	99%	yes	100%
Treatment	piperacillin + amikacin		piperacillin + amikacin	

Table 1: Predictions and advice produced by the VAP model for patient A and B. The treatment advice is based on a utility model in which only overall coverage of suspected pathogens was taken into account.

no leucocytosis and the body temperature was not increased either; the patient did not take antipyretic drugs. Hence, there was much more evidence against than for the presence of VAP, but the signs that were present could also be interpreted as anticipating full-blown pneumonia. The predictions made and the therapy suggested by the model are shown in Table 1. Note that colonisation by the most likely pathogen, Haemophilus influenzae, was later confirmed by the laboratory from sputum cultures.

Patient B had already been mechanically ventilated for 4 days, which was also started immediately after being admitted to the ICU. There were clear symptoms and sign indicating that pneumonia was present, like purulent sputum, leucocytosis, and radiological signs possibly related to pneumonia, although body temperature was again normal. The predictions and advice produced by the model for patient B are also shown in Table 1. Again colonisation by the most likely pathogen, Haemophilus influenzae, was confirmed by sputum cultures, but colonisation by Streptococcus pneumoniae, the pathogen predicted as the second possibility, was not. However, colonisation by Staphylococcus aureus was confirmed.

For these two patients, the system selected the therapy that best covered the suspected pathogens, which was identical in both cases. Although this treatment was able to cover the suspected pathogens with more than 90% chance, the infectious-disease experts viewed the antimicrobial spectrum of the therapy as being far too broad. In fact, this appeared to be a systematic problem encountered in the results for the other patients examined as well. A number of new utility models, in which antimicrobial spectrum was taken into account, were designed, and the tests were repeated. Although these new models did change the ranking of the antibiotics, the results were again not entirely satisfactory.

The results of this preliminary evaluation made clear that the need of the model to cover all possible pathogens tends to result in the prescription of antibiotics with a spectrum that is often too broad, even when this is counterbalanced by a utility model in which the prescription of broad-spectrum antibiotics is discouraged. As is depicted in Figure 11, there are many unlikely pathogens with which a patient might be colonised; the probabilistic model takes all these into account, although in a weighted fashion. In the next section, we shall discuss the results obtained by not only taking the antimicrobial spectrum, but also a probability threshold into account, in an attempt to improve on the results discussed here.

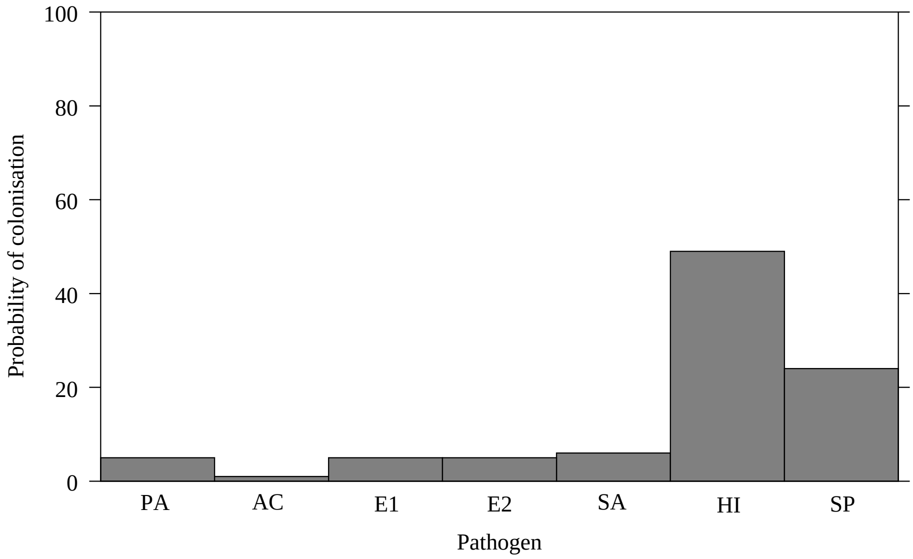

Figure 11: Predicted colonisation profile for patient A.

Another problem that was observed was that taking only the susceptibility of pathogens as determined by the microbiology laboratory into account, may give rise to suboptimal advice. Appropriate antimicrobial therapy recommendations should not only be based on in vitro data, but also on data of the efficacy of the agents at the site of infection in the patient (in vivo data). Unfortunately, whereas the in vitro susceptibility patterns are based on reliable laboratory data, in vivo susceptibility pattern information is heuristic in nature, and hard to establish reliably.

6.2 Study of two utility models, combined with a probability threshold

Data of 38 patients, suspected of having had VAP, and who had been admitted to the ICU in the years 1997 and 1998, were selected from the clinical information system at the ICU. Four different methods, all using the same probabilistic network, were used to determine the optimal antibiotic treatment in this study:

(1) treatment selection using no probability threshold with utility model I (taking into account coverage and five spectrum groups);

(2) treatment selection using a 30% probability threshold, combined with utility model I;

(3) treatment selection using no probability threshold, combined with utility model II (taking into account coverage and eight spectrum groups);

(4) treatment selection using a 30% probability threshold, combined with utility model II.

The infectious-disease specialist who was involved in the design of the model was requested to assess the best three treatment options generated by the system in terms of:

- acceptability (1st–3rd choice, acceptable, unacceptable), and
- antimicrobial spectrum.

The results are shown in the Tables 2, 3, 4 and 5. The following general conclusions can be drawn from these results. The use of a probability threshold had in general a positive effect on the spectrum of the antibiotics chosen by the system: there was a shift of antibiotics into the direction of a more narrow spectrum, although sometimes the spectrum became too narrow (See Tables 2 and 3). A similar effect of the use of a probability threshold could be observed for the ranking of the antibiotics, although the effect was less obvious (See Tables 4 and 5).

Distinguishing more spectrum groups in defining a utility function had an effect quite similar to the probability threshold: the spectrum of the antibiotics selected became more narrow, and there was a shift towards first-choice treatment. However, at the same time the spectrum of selected antibiotics became more often too narrow, and more selected antibiotics were classified as being unacceptable. The combination of utility model I with a probability threshold of 30% seems to offer a good compromise. However, since these results are based on the judgments of one expert, further research has to be performed to investigate whether these judgments are consistent with those of other infectious-disease specialists.

7 Discussion

We have described above the development of a decision-theoretic expert system that is able to assist clinicians in the clinical management of ventilator-associated pneumonia. A number of other decision-support systems have been developed in the past, also using the framework of probabilistic networks (e.g. [2, 22, 33]), sometimes in combination with decision theory (e.g. [18, 19, 28]). The probabilistic-network formalism is one of few formalisms for knowledge representation in which qualitative and quantitative medical knowledge can be easily integrated. It is also a formalism that matches decision theory. Hence, the formalism is eminently suitable for building decision-support systems in domains where uncertainty is of major concern, and where expert knowledge is available to compensate for lack of data. These were exactly the reasons why we adopted the probabilistic-network formalism as a foundation for our research.

There are a number of other researchers who have also studied the problem of computer-based decision support of the clinical management of infectious disease. MYCIN was one of the first systems dealing with this issue, although the system was actually restricted to the clinical areas of sepsis and meningitis [9, 40]. In a preliminary evaluation of the system, it

5 spectrum groups, no threshold				
	1st–3rd		1st choice	
Spectrum	n	%	n	%
$\gg$	36	32	17	45
$>$	33	29	18	47
$=$	31	27	3	8
$\approx$	10	9	0	0
$<$	4	3	0	0
Total	114	100	38	100

5 spectrum groups, 30% threshold				
	1st–3rd		1st choice	
Spectrum	n	%	n	%
$\gg$	10	9	0	0
$>$	25	22	1	3
$=$	53	46	32	84
$\approx$	20	18	1	3
$<$	6	5	4	10
Total	114	100	38	100

Table 2: Assessment of spectrum results for utility model I, without and with probability threshold ($\gg$: much too broad; $>$: too broad; $=$: good; $\approx$: similar; $<$: too narrow).

8 spectrum groups, no threshold				
	1st–3rd		1st choice	
Spectrum	n	%	n	%
$\gg$	12	10	0	0
$>$	9	8	3	8
$=$	44	39	17	45
$\approx$	33	29	5	13
$<$	16	14	13	34
Total	114	100	38	100

8 spectrum groups, 30% threshold				
	1st–3rd		1st choice	
Spectrum	n	%	n	%
$\gg$	3	3	0	0
$>$	2	2	1	3
$=$	75	66	34	89
$\approx$	12	10	0	0
$<$	22	19	3	8
Total	114	100	38	100

Table 3: Assessment of spectrum results for utility model II, without and with probability threshold ($\gg$: much too broad; $>$: too broad; $=$: good; $\approx$: similar; $<$: too narrow).

5 spectrum groups, no threshold				
	1st–3rd		1st choice	
Expert	n	%	n	%
1st	11	10	3	8
2nd	18	16	0	0
3rd	2	2	1	3
acc	79	69	34	89
unacc	4	3	0	0
Total	114	100	38	100

5 spectrum groups, 30% threshold				
	1st–3rd		1st choice	
Expert	n	%	n	%
1st	15	13	4	10
2nd	24	21	22	58
3rd	3	3	1	3
acc	67	59	8	21
unacc	5	4	3	8
Total	114	100	38	100

Table 4: Assessment of ranking for utility model I, without and with probability threshold (acc: acceptable; unacc: unacceptable).

8 spectrum groups, no threshold				
	1st–3rd		1st choice	
Expert	n	%	n	%
1st	11	10	8	21
2nd	20	17	6	16
3rd	1	1	0	0
acc	66	58	11	29
unacc	16	14	13	34
Total	114	100	38	100

8 spectrum groups, 30% threshold				
	1st–3rd		1st choice	
Expert	n	%	n	%
1st	22	19	4	11
2nd	30	26	29	76
3rd	2	2	1	3
acc	39	34	2	5
unacc	21	19	2	5
Total	114	100	38	100

Table 5: Assessment of ranking for utility model II, without and with probability threshold (acc: acceptable; unacc: unacceptable).

was shown that the system performed at least as well as infectious-disease specialists [46]. Rule-based classification systems have been shown to perform well on many occasions (e.g. [1, 14, 30, 29]), but a disadvantage of the rule-based classification approach is that qualitative, causal knowledge, which is readily available in medicine, cannot be employed for their construction. Furthermore, quantitative, probabilistic information cannot be easily integrated with the qualitative knowledge that is available in rules. These two limitations do not exist for probabilistic networks.

The methods that were adopted by Andreassen et al. [3, 4], Leibovici et al. [24], Evans et al. [13], and Warner at al. [44] are related to our work, in the sense that their models were at least partially based on probability theory. The models described by Andreassen at al. [3, 4] are closely related to ours, although they use logistic regression equations to predict the likelihood of the presence of particular pathogens, and the infectious-disease domains are different. Our approach is novel in the sense that it is the first model that uses an explicit representation of the process of bacterial colonisation as a kernel for decision making, which seems in particular appropriate when dealing with hospital-acquired infections. As far as we know, it is also the first system that focusses on the clinical management of the important problem of ventilator-associated pneumonia.

In the near future, the model will be refined by taking the in vivo effects of antimicrobial agents into account, as well as the financial costs and side-effects of antibiotics, to obtain a system that balances different costs and benefits of antibiotic drugs to reach optimal treatment. We intend to embed the resulting system in the clinical information system of the ICU. As Evans et al. have shown, the potential benefits of such a decisions-support system, both for the patient and society, can be significant [13].

References

[1] Adlassnig K-P, Scheithauer W. Performance evaluation of medical expert systems using ROC curves. *Computers and Biomedical Research* 1989; 22: 297–313.

[2] Andreassen S, Woldbye M, Falck B, Andersen SK. MUNIN — A causal probabilistic network for interpretation of electromyographic findings. In: McDermott, J, editor. *Proceedings of the 10th International Joint Conference on Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann, 1987: 366–372.

[3] Andreassen S, Riekehr C, Kristensen B, Schønheyder HC, Leibovici L. Using probabilistic and decision-theoretic methods in treatment and prognosis modeling. *Artif Intell Med* 1999; 15: 121–134.

[4] Andreassen S, Leibovici L, Schønheyder HC, Kristensen B, Riekehr C, Geil Kjaer A, Olesen KG. A decision-theoretic aproach to empirical treatment of bacteraemia originating from the urinary tract. In: Horn W, Shahar Y, Lindberg G, Andreassen S, Wyatt J, editors. *Artificial Intelligence in Medicine*, Lecture Notes in Artificial Intelligence 1999; 1620: 197–206.

[5] Ball MJ, Collen MF, editors. *Aspects of the Computer-based Patient Record*. New York: Springer-Verlag, 1992.

[6] Bartlett JG. *Management of Respiratory Tract Infections*. Baltimore: Williams & Wilkins, 1997.

[7] Bonten MM, Bergmans CJJ, Ambergen AW, De Leeuw PW, Van der Geest S, Stobberingh EE, Gaillard CA. Factors for pneumonia and colonization of respiratory tract stomach in mechanically-ventilated ICU patients. *Am J Respir Crit Care Med* 1996; 154: 1339–1346.

[8] Boutilier C, Dean T, Hans S. Decision-theoretic planning: structural assumptions and computational leverage. *J Artif Intell Re* 1999; 11: 1–94.

[9] Buchanan BG, Shortliffe EH, editors. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* Reading, Massachusetts: Addison-Wesley, 1984.

[10] Cooper GF. A method for using belief networks as influence diagrams. In: *Proceedings of the 4th Workshop on Uncertainty in Artificial Intelligence* 1988: 55–63.

[11] Cotran RS, Kumar V, Robbins SL, Schoen FJ. *Robbins Pathologic Basis of Disease*, 5th ed. Philadelphia: W.B. Saunders, 1994.

[12] Evans RS, Larsen RA, Burke JP, Gardner RM, Meier FA, Jacobson JA, Conti MT, Jacobson JT, Hulse RK. Computer surveillance of hospital-acquired infections and antibiotic use. *JAMA* 1986; 256(8): 1007–1011.

[13] Evans RS, Pestotnik SL, Classen DC, Clemmer TP, Weaver LK, Orme JF, Lloyd JF, Burke JP. A computer-assisted management program for antibiotics and other antiinfective agents. *N Engl J Med* 1998; 338(4): 232–238.

[14] François P, Cremilleux B, Robert C, and Demongeot J. MENINGE: a medical consulting system for child's meningitis: study on a series of consecutive cases. *Artif Intell Med* 1992; 4: 281–292.

[15] Gaag, LC van der. Bayesian belief networks: odds and ends. *The Computer Journal* 1996; 39: 79–113.

[16] Goodman R. *Introduction to Stochastic Models.* Menlo Park, CA: The Benjamin/Cummings, 1988.

[17] Heckerman D, Breese JS. Causal independence for probabilistic assessment and inference using Bayesian networks. *IEEE Transactions on Systems, Man and Cybernetics*, 1996.

[18] Heckerman DE, Horvitz EJ, Nathwani BN. Towards normative expert systems: part I – The Pathfinder project. *Methods Inf Med* 1992; 31: 90–105.

[19] Heckerman DE, Nathwani BN. Towards normative expert systems: part II – probability-based representations for efficient knowledge acquisition and inference. *Methods Inf Med* 1992; 31: 106–116.

[20] Jensen, FV. *An Introduction to Bayesian Networks.* London: UCL Press, 1996.

[21] Kollef MH. The prevention of ventilator-associated pneumonia. *N Engl J Med* 340; 1999: 627–634.

[22] Korver M, Lucas PJF. Converting a rule-based expert system into a belief network. *Med Inf* 1993; 18(3): 219–241.

[23] Lauritzen SL,Spiegelhalter DJ. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society (Series B)* 1987; 50: 157–224.

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

[25] Lucas PJF. Knowledge acquisition for decision-theoretic expert systems. *AISB Quaterly* 1996; **94**: 23–33.

[26] Lucas PJF. Certainty-factor-like structures in Bayesian networks. In: *Proceedings of AIIA'99*, Lecture Notes in Artificial Intelligence. Heidelberg: Springer-Verlag, 2000.

[27] Lucas PJF, Boot H, Taal BG. A decision-theoretic network approach to treatment management and prognosis. *Knowledge-based Systems* 1998; 11: 321–330.

[28] Lucas PJF, Boot H, Taal BG. Computer-based decision-support in the management of primary gastric non-Hodgkin lymphoma. *Methods Inf Med* 1998; 37: 206–219.

[29] Lucas PJF, Janssens AR. Second evaluation of HEPAR, an expert system for the diagnosis of disorders of the liver and biliary tract. *Liver* 1991; 11: 340–346.

[30] Lucas PJF, Segaar RW, and Janssens AR. HEPAR: an expert system for the diagnosis of disorders of the liver and biliary tract. *Liver* 1989; 9; 266–275.

[31] Lucas PJF, Van der Gaag LC. *Principles of Expert Systems*. Wokingham: Addison-Wesley, 1991.

[32] Madoff LC, Kasper DL. Introduction to infectious disease: host–parasite interaction. In: Isselbacher KJ, et al, editors. *Harrison's Principles of Internal Medicine*, 13th edition. New-York: McGraw-Hill, 1994: 485–489.

[33] Middleton B, Shwe MA, Heckerman DE, Henrion M, Horvitz EJ, Lehmann HP, Cooper GF. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base, II. Evaluation of diagnostic performance. *Methods Inf Med* 1991; 30: 256–267.

[34] McIsaac WJ, Goel V. Effect of an explicit decision-support tool on decisions to prescribe antibiotics for sore throat. *Med Decis Making* 1998; 18(2): 220–228.

[35] Miller PL. The evaluation of artificial intelligence systems in medicine. *Comput Programs Biomed* 1986; 22: 5–11.

[36] Papazian L, Thomas P, Garbe L, Guignon I, Thirion X, Charrel J, Bollet C, Fuentes P and Gouin F. Bronchoscopic or blind sampling techniques for the diagnosis of ventilator-associated pneumonia. *Am J Respir Crit Care Med* 1995; 152: 1982–1991.

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

[38] Puterman ML. *Markov Decision Processes*. New York: John Wiley & Sons, 1994.

[39] Shachter RD. Evaluating influence diagrams. *Operation Research* 1986; 34(6): 871–882.

[40] Shortliffe EH. *Computer-Based Medical Consultations: MYCIN*. New York: Elsevier, 1976.

[41] Smith JQ. *Decision Analysis*. London: Chapman and Hall, 1988.

[42] Sox HC, Blatt MA, Higgins MC, Marton KI. *Medical Decision Making*. Boston: Butterworths, 1988.

[43] Srinivas S, Breese J. *IDEAL: Influence Diagram Evaluation and Analysis in Lisp Documentation and Users Guide*. Palo Alto: Rockwell International Science Center, 1993.

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

[45] Wyatt JC, Spiegelhalter DJ. Evaluating medical expert systems: what to test for and how? *Med Inf* 1990; 15(3): 205–217.

[46] Yu VL, Fagan LM, Wraith SM, Clancey WJ, Scott AC, Hannigan J, Blum RL, Buchanan BG, Cohen SN. Antimicrobial selection by a computer. *JAMA* 242; 1979: 1279–1282.