Improving the Therapeutic Performance of a Medical Bayesian Network using Noisy Threshold Models

Stefan Visscher[a], Peter Lucas[b], Marc Bonten[a], and Karin Schurink[a]

[a] University Medical Center Utrecht, Dept. of Internal Medicine and Infectious Diseases, HP F.02.126, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands, {S.Visscher, K.Schurink, M.J.M.Bonten}@umcutrecht.nl

[b] Radboud University Nijmegen, Institute for Computing and Information Sciences, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands, peterl@cs.ru.nl

Abstract. Treatment management in critically ill patients needs to be efficient, as delay in treatment may give rise to deterioration in the patient's condition. Ventilator-associated pneumonia (VAP) occurs in patients who are mechanically ventilated in intensive care units. As it is quite difficult to diagnose and treat VAP, some form of computer-based decision support might be helpful. As diagnosing and treating disorders in medicine involves reasoning with uncertainty, we have used a Bayesian network as our primary tool for building a decision-support system for the clinical management of VAP. The effects of antibiotics on colonisation with various pathogens and subsequent antibiotic choices in case of VAP were modelled in the Bayesian network using the notion of causal independence. In particular, the conditional probability distribution of the random variable that represents the overall coverage of pathogens by antibiotics was modelled in terms of the conjunctive effect of the seven different pathogens, usually referred to as the *noisy-AND gate*. In this paper, we investigate generalisations of the noisy-AND, called *noisy threshold models*. It is shown that they offer a means for further improvement to the performance of the Bayesian network.

1 Introduction

Establishing an accurate diagnosis and choosing appropriate treatment are desirable especially when it concerns critically ill patients. In the intensive care unit (ICU), patients are often severely ill. Patients who depend on respiratory support in the ICU are even more vulnerable than other patients, and are at risk of developing *ventilator-associated pneumonia*, or VAP for short. Thus, it is important to start antimicrobial treatment against VAP as soon as possible in these patients. However, unnecessary antimicrobial treatment will enhance selection of antibiotic-resistant pathogens, which may subsequently cause difficulty in treating future infections adequately. Since only time-consuming and patient-unfriendly diagnostic tests are available for diagnosing VAP, some form

of computer-based decision support could be helpful in the process of early diagnosis and treatment of VAP.

Previously, we have developed a computer-based decision-support system (DSS) that is aimed at assisting physicians in the diagnosis and treatment of VAP. The model underlying the DSS consists of a Bayesian network with an associated decision-theoretic part. The structure as well as the conditional probabilities and utilities were elicitated with the help of two infectious disease specialists. The resulting decision-theoretic model, or influence diagram, was translated into a Bayesian network, and this is the model currently used (Cf. Ref. [1] for details concerning the model and the translation process). The probability of VAP is computed using the diagnostic part of the Bayesian network; the best possible combination of antibiotics can be determined using the therapeutic part of the network.

When prescribing antimicrobial treatment a physician wishes to cover all microorganisms causing the infection, with a spectrum of antibiotics as narrow as possible. This policy aims at preventing the creation of antibiotic resistance and at saving financial costs [2]. This was already taken into account when constructing the DSS, described in more detail in Ref. [1]. To cover as many of the pathogens as possible by the antibiotic treatment advised by the DSS, a noisy-AND gate was used in the Bayesian network for the modelling of the probabilistic interactions of the effects of the prescribed antibiotics on the pathogens. However, it was found that this way of modelling often yields an antibiotic spectrum which is too broad. In the research reported in this paper, noisy threshold functions replace the noisy-AND gate used previously. We investigate whether the therapeutic performance of the Bayesian network for VAP improves in this way. Thus, the aim of the research was to refine the Bayesian network so that it prescribes antibiotics with a spectrum that is less broad.

The paper is organised as follows. In the next section, our earlier work on the development of a Bayesian network that is able to assist physicians in the diagnosis and treatment of VAP is briefly reviewed. In Section 3, the mathematical principles of causal independence models are discussed and noisy threshold functions are introduced. In Section 4, the data and methods used in evaluating the Bayesian networks incorporating the noisy threshold functions are described. The results achieved are commented on in Section 5. The paper is rounded off by some conclusions in Section 6.

2 A Bayesian Network for the Management of VAP

Bayesian networks, or BNs for short, have been introduced in the 1980s as a formalism to compactly represent and reason efficiently with joint probability distributions. Bayesian networks are in particular well suited for the representation of causal relations within a specific domain of expertise.

Formally, a Bayesian network $\mathcal{B} = (G, \Pr)$ is a directed acyclic graph $G = (\mathbf{V}(G), \mathbf{A}(G))$ with set of vertices $\mathbf{V}(G) = \{V_1, \ldots, V_n\}$, corresponding to stochastic variables, here denoted by the same indexed letters, and a set of arcs $\mathbf{A}(G) \subseteq$

$\mathbf{V}(G) \times \mathbf{V}(G)$, representing statistical 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 is factorised respecting the independences 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.

The formalism of BNs supports the kind of the reasoning under uncertainty that is typical for medicine when dealing with diagnosis, treatment selection, planning, and prediction of prognosis. Our medical domain is restricted to patients who are mechanically ventilated and are at risk of developing ventilator-associated pneumonia. Entities that play an important role in the development of VAP and that belong to the diagnostic part of the Bayesian network for VAP include: the duration of *mechanical ventilation*, the amount of *sputum*, *radiological signs*, i.e., whether the chest radiograph shows signs of an infection, body *temperature* of the patient and the number of *leukocytes* (white blood cells) [3]. The structure of the Bayesian network for VAP is shown in Fig. 1. Mechanically ventilated ICU patients become colonised by bacteria. When colonisation of the lower respiratory tract occurs within 2–4 days after intubation, this is usually caused by antibiotic-sensitive bacteria, whereas after one week of intubation often antibiotic-resistant bacteria are involved in colonisation and infection. Such infections are more difficult to treat and immediate start of appropriate treatment is, therefore, important. Duration of hospital stay and severity of illness are associated with an increased risk of colonisation and infection with Gram-negative bacteria. We modelled seven groups of microorganisms, each as one vertex in the Bayesian network. Also, for each modelled microorganism, the pathogenicity, i.e., the influence of that particular microorganism on the development of VAP, was included in the model. The presence of certain bacteria is influenced by antimicrobial therapy. Each microorganism is susceptible to some particular antibiotics. Susceptibility, in this case, is stated as the sensitivity to or degree to which a microorganism is affected by treatment with a specific antibiotic. The susceptibility of each microorganism was taken into account while constructing the model. The infectious-disease experts assigned utilities, by definition quantitative measures of the strength of the preference for an outcome [4], to each combination of microorganism(s) and antimicrobial drug(s) using a decision-theoretic model. These variables are included in the therapeutic part of the Bayesian network for VAP.

3 Causal Independence Modelling

Causal independence is a popular means to facility the specification of conditional probability distributions $\Pr(V_i \mid \pi(V_i))$ involving many parent variables $\pi(V_i)$. Its basic principles and some special forms are briefly discussed below.

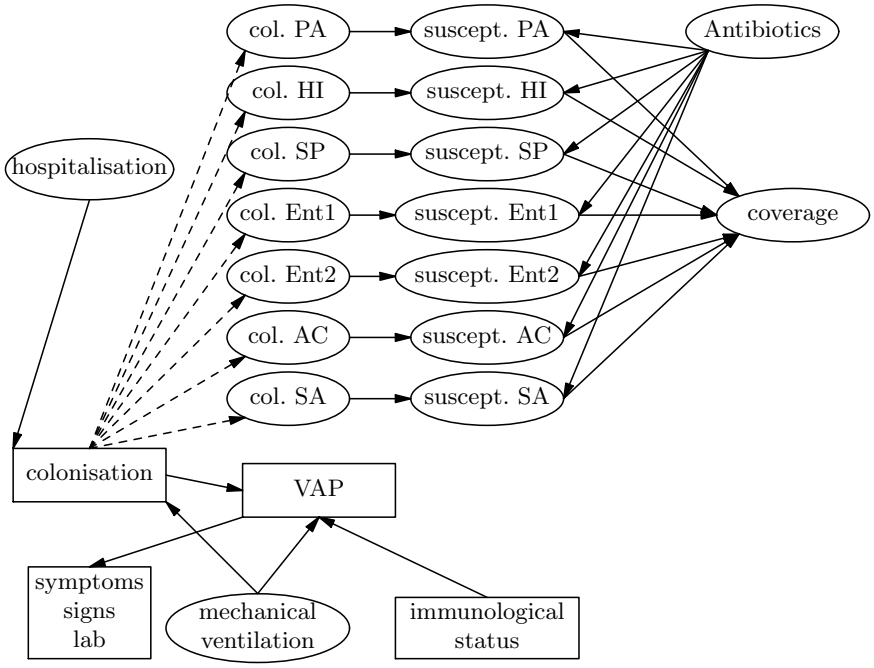

Fig. 1. Abstract model of the Bayesian network for the management of VAP. Colonisation and pneumonia play a central role in this model. The duration of hospitalisation and mechanical ventilation have influence on colonisation (col.) of the patient. PA: Pseudomonas aeruginosa; HI: Haemophilus influenzae; SP: Streptococcus pneumoniae; Ent{1,2}: Enterobacteriaceae{1,2}; SA: Staphylococcus Aureus; AC: Acinetobacter. Each pathogen is susceptible (suscept.) to particular antibiotics and an optimal coverage of the pathogens is what the model tries to achieve. The duration of mechanical ventilation, immunological status and colonisation have influence on the development of VAP. When a patient is diagnosed with VAP, the patient often has symptoms like for example an increased body temperature. Boxes denote entities or processes which are observed; processes that change or can be changed are denoted by ellipses.

3.1 Basic Principles

Consider the conditional probability distribution $\Pr(E \mid C_1, \ldots, C_n)$, where the variable E stands for an *effect*, e.g., coverage, and the variables C_j, $j = 1, \ldots, n$, denote *causes*, e.g., colonisation by pathogens. By taking a number of assumptions into account, which are summarised in Fig. 2, it is possible to simplify the specification of $\Pr(E \mid C_1, \ldots, C_n)$. These assumptions are: (1) the causes C_j are assumed to be mutually independent, and (2) the variable E is conditionally independent of any cause variable C_j given the intermediate variables $I_1, \ldots, I_n$. In our domain the intermediate variable I_j stands for susceptibility of pathogen$_j$

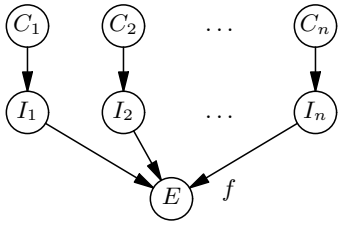

Fig. 2. Causal independence model.

to a specific antibiotic. The, using basic probability theory, it follows that:

$$\Pr(e \mid C_1, \ldots, C_n) = \sum_{I_1,\ldots,I_n} \Pr(e \mid I_1, \ldots, I_n) \prod_{j=1}^{n} \Pr(I_j \mid C_j). \tag{1}$$

Now, if we assume that the probability distribution $\Pr(E \mid I_1, \ldots, I_n)$ that is specified for variable E expresses some deterministic function $f : I_1 \times \cdots \times I_n \rightarrow E$, called an *interaction function*, an alternative formalisation is possible. Using the interaction function f and the causal parameters $\Pr(I_j \mid C_j)$, it follows that [5–7]:

$$\Pr(e \mid C_1, \ldots, C_n) = \sum_{f(I_1,\ldots,I_n)=e} \prod_{j=1}^{n} \Pr(I_j \mid C_j). \tag{2}$$

The result is called a *causal independence model* [5, 8, 6]. In this paper we assume that the function f in Equation (2) is a Boolean function. Systematic analyses of the global probabilistic patterns in causal independence models based on restricted Boolean functions were presented in Ref. [6] and Ref. [9]. However, there are 2^{2^n} different n-ary Boolean functions [10, 11]; thus, the potential number of causal interaction models is huge. However, if we assume that the order of the cause variables does not matter, the Boolean functions become symmetric; formally, an interaction function f is called *symmetric* if

$$f(I_1, \ldots, I_n) = f(I_{j_1}, \ldots, I_{j_n})$$

for any index function $j : \{1, \ldots, n\} \rightarrow \{1, \ldots, n\}$ [11]. The number of different Boolean function reduces then to 2^{n+1}. Examples of symmetric binary Boolean functions include the logical OR, AND, exclusive OR and bi-implication. An example of a general symmetric Boolean functions is the *exact* Boolean function e_k, which is defined as:

$$e_k(I_1, \ldots, I_n) = \begin{cases} \top \text{ if } \sum_{j=1}^{n} \nu(I_j) = k \\ \bot \text{ otherwise} \end{cases} \tag{3}$$

with $k \in \mathbb{N}$, and

$$\nu(I) = \begin{cases} 1 \text{ if } I = \top \\ 0 \text{ otherwise} \end{cases}$$

where $\top$ stands for 'true', and $\bot$ for 'false'. The interaction among variables modelled by the susceptibility, or coverage variables, as shown in Fig. 1, was modelled by assuming f to be a logical AND. The resulting probabilistic model $\Pr(E \mid C_1, \ldots, C_n)$ is sometimes called the *noisy-AND* or *noisy-AND gate*. The probability distribution of the variable that represents the overall susceptibility (coverage in Fig. 1), models the conjunctive effect of the seven different pathogens. This principle is modelled by a probability distribution $\Pr(E \mid C_1, \ldots, C_n)$ that is defined as in Equation (1) by the noisy-AND, yielding the following equation:

$$\Pr(\text{coverage} \mid \text{Col}_1, \ldots, \text{Col}_n, \text{Antibiotics}) = \prod_{j=1}^{n} \Pr(\text{susceptibility-pathogen}_j \mid \text{Col}_j, \text{Antibiotics}).$$

By adopting this modelling approach, the network attempts to cover all pathogens in choosing appropriate antimicrobial treatment.

Evidence has shown that a patient can be colonised by at most 3 pathogens. Therefore, covering the 7 possible groups of pathogens is simply too much and results most of time in choosing antimicrobial treatment with a spectrum that is too broad. This casts doubts on the appropriateness of the noisy-AND for the modelling of interactions concerning coverage of bacteria by antibiotics.

3.2 Threshold Functions

As argued before, clinicians need to be careful in the prescription of antibiotics as they have a tendency to prescribe antibiotics with a spectrum that is too broad. A symmetric Boolean function that is useful in designing a generalised version of the noisy-AND is the *threshold function* τ_k, which simply checks whether there are at least k trues among its arguments, i.e., $\tau_k(I_1, \ldots, I_n) = \top$ (i.e., true), if $\sum_{j=1}^{n} \nu(I_j) \geq k$ with $\nu(I_j)$ equals 1, if I_j equals $\top$ (true) and 0 otherwise [11]. Note that the noisy-AND gate corresponds to the threshold function τ_k with $k = n$. Hence, the noisy-AND can be taken as one extreme of a spectrum of Boolean functions based on the threshold function.

Using the threshold function τ_k with $k \neq 1, n$, may result in a better model. More intuitively, using the noisy threshold functions the network would only cover for 1 ($k = 1$), i.e. noisy-OR, 2 ($k = 2$), 3 ($k = 3$), 4 ($k = 4$), 5 ($k = 5$) or 6 ($k = 6$) pathogens compared to the noisy-AND gate, where all pathogens, i.e. $k = 7$, are taken into account. In the following we therefore investigate properties of the threshold function, and subsequently study its use in improving the Bayesian network model shown in Fig. 1.

3.3 The Noisy Threshold Model

Symmetric Boolean functions can be decomposed in terms of the exact functions e_k as follows [11]:

$$f(I_1, \ldots, I_n) = \bigvee_{k=0}^{n} e_k(I_1, \ldots, I_n) \wedge \gamma_k \tag{4}$$

where γ_k are Boolean constants only dependent of the function f. Using this result, the conditional probability of the occurrence of the effect E given the causes $C_1, \ldots, C_n$ can be decomposed in terms of probabilities that exactly l amongst the intermediate variables $I_1, \ldots, I_n$ are true, as follows:

$$\Pr(e \mid C_1, \ldots, C_n) = \sum_{\substack{0 \leq l \leq n \\ \gamma_l}} \sum_{e_l(I_1,\ldots,I_n)} \prod_{j=1}^{n} \Pr(I_j \mid C_j). \tag{5}$$

Thus, Equation (5) yields a general formula to compute the probability of the effect in terms of exact functions in any causal independence model where an interaction function f is a symmetric Boolean function.

Let us denote a conditional probability of the effect E given causes $C_1, \ldots, C_n$ in a noisy threshold model with interaction function τ_k as $\Pr_{\tau_k}(e \mid C_1, \ldots, C_n)$. Then, from Equation (5) it follows that:

$$\Pr_{\tau_k}(e \mid C_1, \ldots, C_n) = \sum_{k \leq l \leq n} \sum_{e_l(I_1,\ldots,I_n)} \prod_{j=1}^{n} \Pr(I_j \mid C_j). \tag{6}$$

4 Data and Methods

In our attempt to improve the performance of therapeutic advice provided by the Bayesian network for VAP, the following data and methods were used.

4.1 Data

We used a temporal database with 17710 records, each record representing a period of 24 hours of a mechanically ventilated patient in the intensive care unit. The database contains information of 2233 distinct patients, admitted to the ICU of the University Medical Center Utrecht between 1999 and 2002. For 157 of these 17710 episodes, a VAP was diagnosed according to the judgement of two infectious-disease specialists (IDS). We considered the period from admission to the ICU until discharge from the ICU of the patient as a *time series* $\langle X_t \rangle$, $t = 0, \ldots, n_p$, where $t = n_p$ is the time of discharge of patient p. The time-point at which VAP was diagnosed was denoted by t_p^{VAP}, $t = 0 \leq t_p^{VAP} \leq t = n_p$.

For each patient day, we collected the output of the Bayesian network, i.e., the best possible antimicrobial treatment. As reasoning with the network is time-consuming, certainly when varying therapy advice, we (randomly) selected 6 patients, with a total of 40 patient days, who were diagnosed with VAP.

4.2 Therapy Advice

During the period of seven days from the time-point of diagnosis, the patient is treated with antibiotics. Table 1 shows information for the 6 patients using the original Bayesian network. When the number of days following the day of

the diagnosis of VAP is less than 7, we assume that this patient recovered, or died. We furthermore assume that when a patient is colonised by one or more microorganisms on a given day t_c, that after three days, i.e. $t_c+1,\ t_c+2,\ t_c+3$, this patient is still colonised.

The columns in Table 1 have the following meaning:

1. the patient number;
2. the day at which VAP was diagnosed (indicated by prefix '1' instead of '0');
3. the number of days a patient has been mechanically ventilated;
4. the microorganism(s) found in the sputum culture; abbreviations:
   - acinetb = Acinetobacter;
   - entbct{1,2} = Enterobacteriaceae{1,2};
   - hinflu = H. influenzae;
   - paeru = P. aeruginosa;
   - spneumon = S. pneumoniae.
5. antibiotics, as mentioned above, can be divided in spectral groups. Used abbreviations are:
   - v : very narrow;
   - n : narrow;
   - i : intermediate;
   - b : broad.

Comparison of antimicrobial spectrum imposes some difficulty. There are, for example, several intermediate-spectrum antibiotics available, possibly produced by different vendors. Thus, it is possible that in the table two different antibiotics are mentioned, even through they have the same effect. For example, in the table we see that for patient 2 on day 10, the ICU physician prescribes *cefpirom*, whereas the Bayesian network advises to prescribe *ceftazidime*. The antimicrobial therapy prescribed by the ICU physician and corresponding spectrum indicated between parentheses, as well as the therapy advice given by the Bayesian Network with associated spectrum, are mentioned in the last two columns of the table. Note that 'none' means that the Bayesian network advises not to prescribe any antibiotics.

4.3 Methods

In order to improve the therapeutic performance of the Bayesian network, the network was inspected in detail. Points for possible improvement that were identified included the utilities, used in the selection of antibiotics, and the noisy-AND, used as a basis for the assessment of the conditional probability distribution

$$\Pr(\text{Coverage} \mid \text{Col}_1, \ldots, \text{Col}_7, \text{Antibiotics}).$$

Based on the inspection, the following actions were taken:

Table 1. This table shows the therapeutic advises by the ICU physician (column 5) and the Bayesian network (column 6), as well as the clinical culture data (column 4): acinetb = Acinetobacter; entbct{1,2} = Enterobacteriaceae{1,2}; hinflu = H. influenzae; paeru = P. aeruginosa; spneumon = S. pneumoniae; negative = 'no microorganisms found'.

Patient	VAP	day	Colonised by	Antibiotics selected by Physician	BN
1	1	6	entbct1	augmentin (n)	meropenem (b)
1	0	7		cefpirom (i)	meropenem (b)
1	0	8		cefpirom (i)	meropenem (b)
2	1	7	paeru, hinflu, entbct2	cefpirom (i)	clindam.+ciprox. (b)
2	0	8		cefpirom (i)	clindam.+ciprox. (b)
2	0	9		cefpirom (i)	clindam.+ciprox. (b)
2	0	10	paeru	cefpirom (i)	ceftazidime (i)
2	0	11		ciproxin (b)	ceftazidime (i)
2	0	12		ciproxin (b)	ceftazidime (i)
2	0	13		ciproxin (b)	ceftazidime (i)
2	0	14		ciproxin (b)	none
3	1	5		augm/erytro/gent (i)	meropenem (b)
3	0	6	entbct1	augm/erytro (i)	meropenem (b)
3	0	7		erytro/ceftriaxon (i)	meropenem (b)
3	0	8	entbct1	erytro/ceftriaxon (i)	meropenem (b)
3	0	9		ceftriaxon (i)	meropenem (b)
3	0	10		ceftriaxon (i)	meropenem (b)
3	0	11	negative	ceftriaxon (i)	ceftriaxone (i)
3	0	12		ceftriaxon (i)	ceftriaxone (i)
4	1	30	acinetb, entbct1	cefpirom (i)	meropenem (b)
4	0	31		cefpirom (i)	meropenem (b)
4	0	32	acinetb	cefpirom (i)	meropenem (b)
4	0	33	acinetb	cefpirom (i)	meropenem (b)
4	0	34		cefpirom (i)	meropenem (b)
4	0	35		cefpirom (i)	meropenem (b)
4	0	36		cefpirom (i)	meropenem (b)
4	0	37		cefpirom (i)	meropenem (b)
5	1	5	hinflu, spneumon	augmentin (n)	meropenem (b)
5	0	6		augmentin (n)	meropenem (b)
5	0	7		augmentin (n)	meropenem (b)
5	0	8		augmentin (n)	meropenem (b)
5	0	9		augmentin (n)	none
6	1	11	paeru, entbct1	augm/pipcil (i)	clindam.+ciprox. (b)
6	0	12		pipcil (i)	clindam.+ciprox.(b)
6	0	13		pipcil (i)	clindam.+ciprox. (b)
6	0	14		pipcil (i)	clindam.+ciprox. (b)
6	0	15		pipcil (i)	clindam.+ciprox. (b)
6	0	16		pipcil/ciprox.(b)	clindam.+ciprox. (b)
6	0	17	paeru, entbct1	ciproxin (b)	clindam.+ciprox. (b)
6	0	18		ciproxin (b)	clindam.+ciprox. (b)

Table 2. Old and new utilities.

	Utilities	
Spectrum	**Old**	**New**
none	29	29
very narrow	96	96
narrow	89	89
intermediate	82	82
broad	**71**	**60**

- It became clear that the preferences of antibiotics with a broad spectrum were overestimated by the experts while assigning utilities. By giving the broad-spectrum antibiotics a lower utility, it was expected that this might result in a more appropriate treatment advice. On a scale between 0 and 100, with 0 representing 'not preferred' and 100 'preferred' new utilities were assessed by infectious disease experts; both the old and new utilities are summarised in Table 2. In redefining the utilities, it was assumed that each patient had VAP and that we wished to cover all pathogens present.
- We studied the use of noisy threshold functions as explained in Section 3.2, in achieving a better therapeutic performance by the Bayesian network.

Two infectious-disease specialists were requested to prescribe antibiotics for each of the 6 patients. Their treatments were considered to be the *gold standard*. This allowed us to validate the outcomes of the study.

5 Results

When prescribing antibiotics, choice of the spectrum should be based on taking into account the susceptibilities of causative pathogens. When the number of different causative pathogens increases, the necessary coverage will become more broad and often more than one antibiotic will be prescribed. Preliminary results, shown in Table 1, indicate that the antimicrobial spectrum advised by the Bayesian network is often broad, even when only two causative pathogens are present. This effect is summarised in Table 3 in the column with header 'Noisy-AND old'. Note that the treatments given by the ICU doctors included in the table are not considered to be *gold standard*, as is it known that ICU doctors have a tendency to prescribe antibiotics with a spectrum that is often too broad.

Table 3 indicates in the column with header 'Noisy-AND new' that the redefinition of the utilities already resulted in a better therapeutic performance. Yet, the Bayesian network still advised antibiotics with a spectrum that was often too broad. Hence, there was a clear need for further refinement of the network, which was subsequently undertaken using noisy threshold models.

Table 3. Results of the prescription of antibiotics to 6 patients with VAP according to: ICU physician, the original network (noisy-AND old), network with new utilities, network with noisy threshold with $k = 3$, $k = 4$ and $k = 5$ in comparison to the Infectious Disease Specialists (IDS). Abbreviations of antibiotic spectrum: o: none; v: very narrow; n: narrow; i: intermediate; b: broad.

Patient / **Model**	**1**	**2**	**3**	**4**	**5**	**6**	**total**
ICU Physician	1n 2i	4i 4b	8i	8i	5n	5i 3b	6n 27i 7b
Noisy-AND old ($k = 7$)	3b	1o 4i 3b	2i 6b	8b	1o 4b	8b	2o 9i 29b
Noisy-AND new ($k = 7$)	3i	8i	1n 4i 3b	8b	5i	3i 5b	1n 23i 16b
Noisy threshold ($k = 3$)	3v	8v	8v	8v	5v	8v	40v
Noisy threshold ($k = 4$)	3v	4v 4n	4v 4n	2v 4n 2i	1v 4n	8n	14v 24n 2i
Noisy threshold ($k = 5$)	1n 2i	4n 4i	5v 1n 2i	2n 6i	3n 2i	3n 5i	5v 14n 21i
IDS *gold standard*	3n	8i	8i	8i	5n	8i	8n 32i

For each noisy threshold model, we collected the output in the same manner as for Table 1. We have summarised the resulting antibiotic spectra per patient in Table 3 for thresholds $k = 3$, $k = 4$, $k = 5$ and the two noisy-ANDs. The performance for the noisy-OR model (threshold function with $k = 1$), $k = 2$ (both not in the table), and $k = 3$ were rather poor: for all patients the resulting antimicrobial spectrum was too narrow. The networks with threshold function with $k = 6$ (not in the table) and the original network with the noisy-AND (threshold function with $k = 7$), had a poor therapeutic performance as well, but here the prescribed antibiotic spectrum was too broad in most of the cases. The Bayesian networks with threshold functions with $k = 4$ and $k = 5$, however, performed relatively well.

6 Conclusions and Discussion

In this paper, we have shown that by reconsidering the modelling of interactions between variables in a Bayesian network, it is possible to improve its performance. We used a Bayesian network for the diagnosis and treatment of ventilator-associated pneumonia as an example. Intensive use was made of the theory of causal independence, which not only facilitates the assessment of probability tables by allowing specifying a table in terms of a linear number of parameters $\Pr(I_j \mid C_j)$, but also allows taking into account domain characteristics [6]. This was clearly shown for our Bayesian network concerning VAP, where motivation was derived from the domain of infectious disease, indicating that a noisy threshold model might be appropriate for the modelling of the interaction between pathogens and antimicrobial susceptibility. It appeared that a noisy threshold function with $k = 5$ yielded the best results, according to the gold standard, i.e., the infectious disease experts. This also provides evidence that the noisy OR and noisy AND, which are very popular in Bayesian network modelling, might not be the best interaction functions for other application areas as

well. In the near future, we intend to test our findings on a larger sample of our database.

To conclude, it was shown that the noisy threshold model is useful from a practical point of view by using it as a basis for the refinement of an existing real-world Bayesian network for the management of critically ill patients.

Acknowledgements This research, undertaken in the TimeBayes project, was funded by the Netherlands Science Foundation (NWO) under the ToKeN2000 programme. We would like to thank Rasa Jurgelenaite for her help with the theoretical work concerning noisy threshold models.

References

1. Lucas PJF, Bruijn NC de, Schurink K, and Hoepelman A. A probabilistic and decision-theoretic approach to the management of infectious disease at the icu. *Artificial Intelligence in Medicine*, 19(3):251–279, 2000.
2. Bonten MJM. Prevention of infection in the intensive care unit. *Current Opinion in Critical Care*, 10(5), 2004.
3. Bonten MJM, Kollef MH, and Hall JB. Risk factors for ventilator-associated pneumonia: from epidemiology to patient management. *Clinical Infectious Diseases*, 38(8), 2004.
4. Sox HC, Blatt MA, Higgins MC, and Marton KI. *Medical Decision Making*. Butterworth-Heinemann, 1988.
5. Heckerman D. Causal independence for knowledge acquisition and inference. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*, 1993.
6. Lucas PJF. Bayesian network modelling through qualitative patterns. *Artificial Intelligence*, 163:233–263, 2005.
7. Zhang LH and Poole D. Exploiting causal independence in bayesian networks inference. *Journal of Artificial Intelligence Research*, 5:301–328, 1996.
8. Heckerman D and Breese JS. Causal independence for probabilistic assessment and inference using bayesian networks. *IEEE Transactions on Systems, Man and Cybernetics*, 26(6), 1996.
9. Jurgelenaite R and Lucas PJF. Exploiting causal independence in large bayesian networks. *Knowledge Based Systems Journal*, 18(4–5), 2005.
10. Enderton HB. *A Mathematical Introduction to Logic*. Academic Press, 1972.
11. Wegener I. *The Complexity of Boolean Functions*. John Wiley & Sons, 1987.