APPLICATION OF THE CLOSED-WORLD ASSUMPTION TO A BAYESIAN NETWORK FOR OPTIMAL TREATMENT SELECTION

S. Visscher*, P.J.F. Lucas**, C.A.M. Schurink* and M.J.M. Bonten*

* University Medical Center Utrecht, Department of Medicine, Division of Acute Medicine & Infectious Diseases, HP F.02.126, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands
** Radboud University Nijmegen, Institute for Computing and Information Sciences, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands

S.Visscher@umcutrecht.nl

Abstract: Managing disorders in critical care is a challenging task, as when a disorder is missed and thus treatment is delayed this may be the cause of death of a patient, as many of the patients in critical care are severely ill. Ventilator-associated pneumonia occurs in patients who are mechanically ventilated in intensive care units. As diagnosing and treating infections involves reasoning with uncertainty, we used a Bayesian network as our primary tool for building a decision-support system. Prescribing antimicrobial treatment is a trade-off between minimising antimicrobial spectrum and maximising coverage of pathogens to be treated. To support clinicians in performing their task, the Bayesian network computes optimal treatment for a patient given colonisation data. We analysed the therapeutic performance of the Bayesian network by comparing its output to the gold standard, i.e., antibiotic treatment prescribed by two infectious-disease specialists. It is shown that by using the concept of the Closed-World Assumption with respect to colonisation data it was possible to achieve accurate results.

Introduction

Patients who are admitted to the Intensive Care Unit (ICU) are often severely ill and are prone to colonisation by hospital-acquired (nosocomial) bacteria. Some of these bacteria are more pathogenic[1] than others. Our medical domain is restricted to mechanically ventilated patients who are at risk of developing an infection of the lower respiratory tract. This infection is called ventilator-associated pneumonia, or VAP for short. To treat pathogens that have caused an infection, antimicrobial treatment is needed.

Prescribing optimal treatment is a difficult task for physicians. They have the tendency to prescribe broad spectrum antibiotics to cover as much pathogens as possible. However, this stimulates the development of resistance of pathogens to specific antibiotics, which will eventually reduce the effectiveness of these antibiotics in time [1]. Therefore, more and more research is being performed in the area of standardisation of antibiotic prescription; guidelines and protocols are well-known practical outcomes of such efforts. In our research, we focus on supporting physicians in prescribing antimicrobial treatment for ICU patients by means of a Bayesian network (BN) augmented by a decision-theoretic model [2]. Bayesian networks 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 uncertain causal relations within a specific domain of expertise.

This paper describes our attempt to improve the therapeutic performance of our Bayesian network. We used real patient data to test the accuracy of the network in selecting optimal antimicrobial treatment. It is shown that performance improvement can be achieved by using the concept of the Closed-World Assumption (CWA) from database theory and logic [3].

The paper is organised as follows: Bayesian networks and the Bayesian network for the treatment of VAP we have developed previously are first described, followed by a brief description of the basic ideas underlying the CWA. The exploitation of the CWA in the context of probability theory is studied next. The practical usefulness of these ideas are subsequently investigated for the Bayesian network concerning VAP. The paper is rounded off by some conclusions.

Materials and Methods

A Bayesian network for VAP

Formally, a *Bayesian network*, BN for short, is defined as a pair $\mathscr{B} = (G, \Pr)$, where $G = (\mathbf{V}(G), \mathbf{A}(G))$ is a directed acyclic graph with a set of vertices $\mathbf{V}(G) = \{V_1, \ldots, V_n\}$, corresponding one to one to stochastic variables, here denoted by the same indexed letters, and a set of arcs $\mathbf{A}(G) \subseteq \mathbf{V}(G) \times \mathbf{V}(G)$, and Pr is a joint probability distribution $\Pr(V_1, \ldots, V_n)$ representing statistical dependences and independences among the variables, respecting the independences represented in the graph, as follows:

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

[1] Pathogenicity is the ability of an organism of causing an individual to get ill.

ISSN: 1727-1983

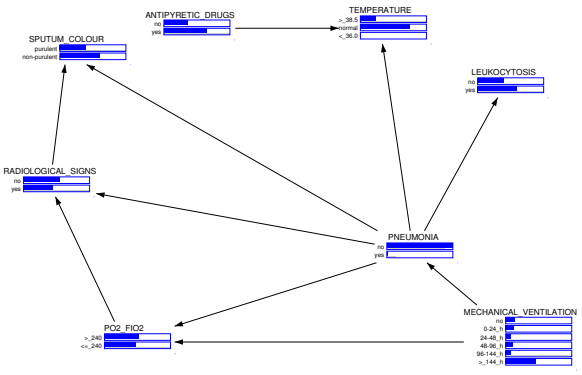

Figure 1: Diagnostic part of the Bayesian network; the 'PNEUMONIA' vertex links this diagnostic part and the therapeutic part together.

where $\pi(V_i)$ stands for the variables corresponding to the parents of vertex V_i. In the following, upper-case letters, such as X, stand for (free) variables, whereas lower-case letters, such as x (short for $X = yes$) or $\neg x$ (short for $X = no$), stand for (collections) of values of variables.

The formalism of BNs supports the kind of reasoning under uncertainty that is typical for medicine when dealing with diagnosis, treatment selection and planning, and prediction of prognosis. A Bayesian-network model concerning VAP was previously constructed with the help of two infectious-disease domain experts, who helped in establishing the network's structure and estimated all conditional probabilities required [2]. This Bayesian network can be seen as consisting of a diagnostic part, modelling signs and symptoms of VAP, and a therapeutic part, that models the temporal evolution of the disease based on duration of stay in the ICU and duration of mechanical ventilation. These parts are linked together through the 'PNEUMONIA' vertex (the 'PNEUMONIA' vertex has one incoming arc from the therapeutic part).

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) [4]. Each of these entities is modelled as one vertex in the diagnostic part of the Bayesian network for VAP, as shown in Figure 1.

The *therapeutic part* of the network models the situation of a patient from the colonisation and development of pneumonia as temporal processes, to the selection of optimal antimicrobial treatment. We have modelled seven groups of microorganisms, each as one vertex in the Bayesian network. These microorganisms are: Pseudomonas aeruginosa; Haemophilus influenzae; Streptococcus pneumoniae; two groups of Enterobacteriaceae, depending on which antibiotics these are susceptible to; Staphylococcus aureus and Acinetobacter spp. Also, for each modelled microorganism the pathogenicity was included in the model; the pathogenicities were assumed to be equal for each microorganism.

The presence of certain bacteria is influenced by antimicrobial therapy. Each microorganism is susceptible[2] to some particular antibiotics, and these susceptibilities were taken into account while constructing the model. The infectious-disease experts assigned utilities[3] to each combination of microorganism(s) and antimicrobial drug(s) using a decision-theoretic model [5].

Modelling joint interactions

To model the probabilistic interaction of the various pathogens on the likelihood of development of pneumonia and overall susceptibility, the notion of *causal independence* was used [6, 7]. For example, the interaction among susceptibility or coverage variables, as shown in Figure 2, was expressed by a logical-AND gate. The probability distribution of the vertex that represents the overall susceptibility or coverage is expressed as the *conjunctive effect* of the seven different pathogens, which can be defined formally as follows:

$$\Pr(\text{coverage} \mid \text{at}) = \prod_{i=1}^{n} \Pr(\text{susceptibility-pathogen}_i \mid \text{at}) \quad (1)$$

Using the logical-AND gate, the model tries to cover all pathogens, i.e., the probability $\Pr(\text{coverage} \mid \text{at})$ for all possible antimicrobial treatment values 'at' of the variable 'ANTIBIOTICS' is computed. This is balanced against the broadness of the antimicrobial spectrum. Thus, there is a trade-off between coverage and broadness of antimicrobial spectrum of the prescribed treatment.

Preliminary evaluation

Preliminary evaluation of the performance of the current network model learnt us that the model, generally, advised broad spectrum antibiotics, even when the patient was only colonised by one or two pathogens. When prescribing antibiotics, the spectrum should be chosen based on information concerning the susceptibilities of causative pathogens. In general, when increasing the number of different possible causative pathogens, there will be a need for broader coverage and, thus, more different antibiotics. If the patient is only infected by one or two pathogens, prescription of one narrow to intermediate spectrum antibiotic is often sufficient. Our conclusion was that there was room for improvement of the model. More detailed evaluation results that support this claim are discussed below.

The closed-world assumption

In this paper, we study the idea from Artificial Intelligence (AI) concerning the assumption that unknown, absent data can be taken as being negative. We investigate

[2] Susceptibility, here, is stated as the sensitivity to or degree to which a microorganism is affected by treatment with a specific antibiotic.

[3] Utility: by definition a quantitative measure of the strength of the preference for an outcome.

ISSN: 1727-1983

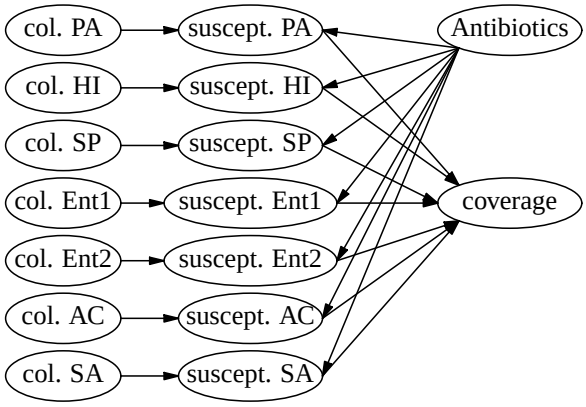

Figure 2: Most important fragment of the therapeutic part of the Bayesian network. PA: Pseudomonas aeruginosa; HI: Haemophilus influenzae; SP: Streptococcus pneumoniae; Ent{1,2}: Enterobacteriaceae{1,2}; AC: Acinetobacter spp.; SA: Staphylococcus aureus. Each pathogen is susceptible (suscept.) to particular antibiotics and an optimal coverage of the pathogens is what the model tries to achieve.

whether this idea can have a positive impact on the performance of a Bayesian network. We first describe the basic ideas from the perspective of AI, followed by how these ideas can be exploited in the context of probability theory.

A typical logical AI system stores its knowledge about its domain as a finite set Γ of first-order logic formulas. To answer queries, the system will have to decide whether or not a formula, say φ, can be obtained by performing logical deductions on Γ; formally: $\Gamma \models \varphi$, meaning that φ is a logical consequence of our knowledge represented in the *knowledge base* Γ. A typical use of this logical approach is in rule-based expert systems, where Γ is a collection of logical rules, and φ could be something such as the diagnosis of a disease based on entered *patient findings* or *evidence* E:

$$\Gamma \cup E \models \varphi$$

Thus, the knowledge base Γ is static, and augmented by non-static, patient-specific findings E to conclude about a disease φ, which is then interpreted as a diagnosis, or prognosis, etc., depending on the medical purpose of the knowledge in Γ.

This model has been proven to be quite useful for various tasks requiring knowledge about a domain, however, it has its limitations. Let us assume that what we know about our world is stored in Γ. When we want to know something about, for example, a certain φ, we search in Γ for information concerning this φ. It might be the case that it is not possible to conclude φ from Γ, indicating that somehow our knowledge base Γ is incomplete. Clearly, this is often the case, as it may be impossible to ensure that a knowledge base offers a complete coverage of all the knowledge in a domain of concern. The assumption that a knowledge base is incomplete is known as the *open world assumption* [3, 8].

However, when the knowledge base Γ cannot tell us anything about φ, we may also assume that the negation of φ, i.e., $\neg\varphi$, holds and can be added to Γ. Formally, we have that if $\Gamma \not\models \varphi$, we assume that $\neg\varphi$ is a member of a set of assumptions or beliefs B, i.e., B is the smallest set of negative assumptions or beliefs, such that for each $\psi \in B : \Gamma \not\models \psi$ and $\Gamma \not\models \neg\psi$, where $\psi \equiv \neg\varphi$. Now, if it is the case that

$$\Gamma \cup B \models \neg\varphi$$

then we say that $\neg\varphi$ is in the CWA-augmented knowledge base Γ, formally $\neg\varphi \in \mathrm{CWA}(\Gamma)$. This is called the *closed world assumption* [3, 8]. It is used to provide a default, negative solution in the absence of a positive solution.

There are at least two situations where the CWA is used. The first is where it is assumed that a knowledge base contains all relevant facts. This is common in corporate databases. That is, the information it contains is assumed to be complete. The second situation is where it is known that the knowledge base is incomplete (does not have enough information to produce an answer to a question) and a decision must be made without complete information — a situation familiar to most people. The closed world assumption is designed to solve a reasoning problems in both of these situations.

In medicine, it is unlikely that all information concerning a patient is known, as only the results of a selected set of tests are known. Thus, it is common to assume that unless explicitly stated, everything not known about a patient is normal. For example, if the blood pressure of a patient is unknown, and there are no indications that absence of information concerning the blood pressure may be a medical mistake, it is normally assumed that the blood pressure is *not* high. Hence, the CWA in medicine is commonly used to interpret data concerning a patient; without it, it is hard to draw clear conclusions. Here we, therefore, assume that the CWA is only used to handle patient data. This means that in reasoning with logical rules in a knowledge base Γ and patient findings E, such that

$$\Gamma \cup E \models \varphi$$

we compute the CWA, not of the knowledge base Γ, but of the set of patient findings E in the context of Γ. Hence, logical reasoning in medicine can be formalised as reasoning according to the following definition

$$\Gamma \cup \mathrm{CWA}_\Gamma(E) \models \psi$$

where if $\neg\varphi \in \mathrm{CWA}_\Gamma(E)$: $\Gamma \cup E \not\models \varphi$, $\Gamma \cup E \not\models \neg\varphi$ and $\psi \equiv \neg\varphi$.

The CWA and probability theory

Despite the fact that the CWA, as described above, is useful in a medical context, a limitation is that, being fully based on logic, it is unable to deal with uncertainty. Thus, there is a need to extend the ideas described above towards probability theory.

In the context of Bayesian networks, the use of the CWA is comparable, but somewhat different. We use as

ISSN: 1727-1983

much information as available from our clinical database to fill in the nodes of the network. The information used for inference in the network is called *evidence*. Let us assumed that we have particular evidence e concerning a patient, excluding a variable X, such that $0 < \Pr(X = yes \mid e) < 1$. This is interpreted as saying that X is not fully known. However, if variable X is potentially part of the patient evidence, then we may, adopting the CWA, assume that $X = no$ holds, thus as a consequence:

$$\Pr(X = yes \mid e \cup \{X = no\}) = 0$$

We have that $\mathrm{CWA}_{\Pr}(e) = e \cup \{X = no\}$. This means that we always compute

$$\Pr(X \mid \mathrm{CWA}_{\Pr}(E))$$

for any set of evidence E.

With regard to our Bayesian network about VAP, we have applied the CWA to data of sputum cultures, but in a restrictive fashion. When no pathogens are found, but a sputum culture has been taken, it is assumed that all pathogens are absent. This is a negative test result and not really an instance of the CWA. When culture data is not missing for a specific patient day, and one or more of the seven possible groups of pathogens are found positive, we deduce that the other groups of pathogens will be absent. These negative culture data will then augment the patient evidence *e*, and this is an instance of the CWA.

Evaluation of therapeutic performance

The performance of the Bayesian network, before and after taking into account the CWA for probability theory, was subsequently evaluated. The following approach was taken:

(1) We first assume that when a patient is colonised by one or more microorganisms on a given day t_c, that on the three successive days, i.e. t_c+1, t_c+2, t_c+3, this patient is still colonised. This assumption seems valid, as (1) clinical cultures usually are not performed daily and (2) when treated with antibiotics, microorganisms are not eradicated immediately.
(2) Secondly, we interpret the clinical culture data as follows:
   (a) When no microorganisms are found, this is interpreted as a *negative* result for the seven nodes in the network we fill in the evidence: 'colonisation_by_pathogen$_i$ = no', $1 \leq i \leq 7$.
   (b) When one or more pathogens are found positive in the culture, these pathogens are processed in the network as being present; for the other pathogens, however, we assume that they are absent. This is in instance of the CWA in probability theory.
   (c) When on a specific day no cultures were performed, i.e. culture data is missing, no evidence is filled in in the network.

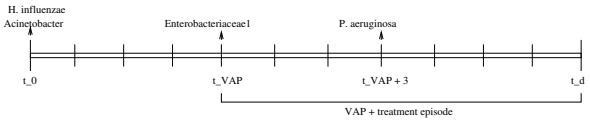

Figure 3: Example time line. On t_0, clinical cultures showed that this patient was colonised with H. influenzae and Acinetobacter spp. Four days later, here denoted by t_{VAP}, the patient was diagnosed with VAP and antimicrobial treatment was started to cover the Enterobacteriaceae1 pathogen and on day 3 of the treatment pathogen P. aeruginosa. At time point t_d, the patient was discharged from the ICU.

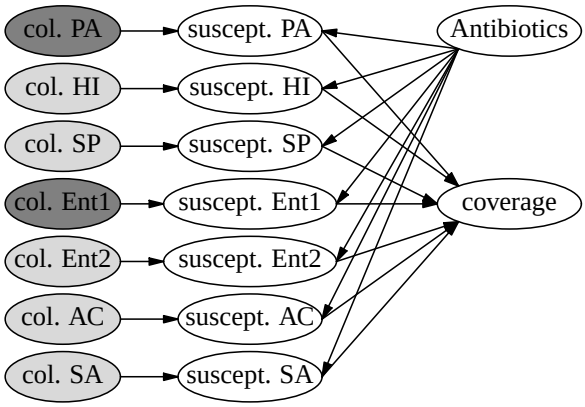

Figure 4: The therapeutic part of the network including clinical evidence as shown on the time line on the third day of treatment ($t_{VAP}+3$). Positive evidence is denoted by the darkest shaded ellipses, whereas the lightest shaded ellipses denote the pathogens which are assumed negative, using the CWA.

To test whether making these assumptions improves the therapeutic performance of the network, we used a temporal database of 17710 records for this purpose. This database contains data for more than 2000 patients, admitted to the ICU between 1999 and 2002 in the University Medical Center Utrecht, The Netherlands. For 157 of these 17710 episodes, a VAP was diagnosed according to the judgment of two infectious-disease specialists (IDS). During the period of seven days from time-point of diagnosis, the patient is treated with antibiotics. When the number of days after the day of VAP is less than 7, we assume that this patient recovered, or died. See Figure 3 for an example of a time line which shows the evidence and actions for a patient from the time point of admission to the ICU until the time point of discharge from the ICU. Using the CWA and the clinical evidence as shown on the time line, on day $t_{VAP}+3$, i.e. on the third day of treating this VAP patient, the network would look as shown in Figure 4.

Results

Two infectious-disease specialists selected optimal antimicrobial treatment for the same patient data as we used

ISSN: 1727-1983

Table 1: Gold standard. P1 and P2 represent the causative pathogens. Abbreviations used for the pathogens: SA: S. aureus; SP: S. pneumoniae; HI: H. influenzae; En1: Enterobacteriaceae1; En2: Enterobacteriaceae2; AB: Acinetobacter; PS: P. aeruginosa. Freq. stands for 'frequency'. Abbreviations for antibiotic spectrum: v = very narrow; n = narrow; i = intermediate; b = broad.

P1	P2	Freq.	Antibiotics	Spectrum
SA		25	Floxapen	vn
SP		4	Penicylin	vn
HI		8	Augmentin	n
En1	En2	4	Ceftriaxone	n
En2	SA	4	Ceftriaxone	n
SA	SP	3	Floxapen	vn
HI	SP	6	Augmentin	n
En1		27	Ceftriaxone	n
En2		18	Ceftriaxone	n
SA	HI	5	Ceftriaxone	n
AB	En1	3	Ceftriaxone	n
PA		19	Ceftazidime	i
En2	AB	3	Ceftriaxone	n
HI	PA	1	Ceftazidime	i
En2	SP	1	Cotrimoxazol	n
PA	En2	5	Ciproxin	b
AB	SA	1	Cotrimoxazol	n
PA	En1	3	Ceftazidime	i
En1	HI	2	Cotrimoxazol	n
PA	AB	1	Ceftazidime	i
SA	PA	1	Tazocin	b
SA	En1	2	Ceftriaxone	n
AB		6	Cotrimoxazol	n
En2	HI	1	Ceftriaxone	n
		4	Cotrimoxazol	n

Table 2: Results. The column 'Old' denotes the results for the original Bayesian network, whereas 'New' gives information about the performance of the Bayesian network, when using the CWA. Abbreviations for antibiotic spectrum: v = very narrow; n = narrow; i = intermediate; b = broad.

	Antibiotics & Spectrum			
Freq.	Old	SP	New (CWA)	SP
25	Clinda+aztr	i	Floxapen	vn
4	Clinda+aztr	i	Peniciline	vn
8	Clinda+aztr	i	Erythromycin	vn
4	Clinda+aztr	i	Clinda+aztr	i
4	Clinda+aztr	i	Clinda+aztr	i
3	Clinda+aztr	i	Floxapen	vn
6	Clinda+aztr	i	Erythromycin	vn
27	Meropenem	b	Ceftriaxone	n
18	Clinda+aztr	i	Clinda+aztr	i
5	Clinda+aztr	i	Ceftriaxone	n
3	Meropenem	b	Meropenem	b
19	Ceftazidime	i	Ceftazidime	i
3	Meropenem	b	Cotrimoxazol	n
1	Ceftazidime	i	Ceftazidime	i
1	Clinda+aztr	i	Clinda+aztr	i
5	Clinda+aztr	i	Clinda+aztr	i
1	Meropenem	b	Cotrimoxazol	n
3	Ceftazidime	i	Ceftazidime	i
2	Meropenem	b	Ceftriaxone	n
1	Meropenem	b	Ceftazidime	i
1	Clinda+aztr	i	Clinda+aztr	i
2	Meropenem	b	Ceftriaxone	n
6	Clinda+aztr	i	Cotrimoxazol	n
1	Meropenem	b	Clinda+aztr	i
4	Clinda+aztr	i	Metronidazole	vn

in our analysis. By doing so, we were able to compare their therapy advice, here considered the *gold standard*, to the treatment selected by the Bayesian network. The information shown in Table 1 can be interpreted as follows: each combination of causative pathogens was put in the first column. The second column denotes the number of occurrences of each combination, adding up to the total number of VAP patients, i.e., 157. The third column is the antibiotic treatment, prescribed by the infectious-disease experts. Broadness of spectrum, matching to the the treatment, can be found in the last column. The results of the evaluation of the Bayesian network, with and without using the CWA, are shown in Table 2.

To be able to draw conclusions from Tables 1 and 2, we have compared them. We have done this in such a way that we are able to say for each group of causative pathogens whether the advised antibiotics by the models are right or wrong. Wrong here means that the antibiotic spectrum, as advised by one of the two interpretations of the Bayesian network, is either too narrow or too broad. Table 3 summarises the results of this comparison.

Conclusions and Discussion

In this paper, we have described a way to improve the performance of a Bayesian network for treatment selection in patients with VAP. The way in which data about sputum cultures was interpreted previously gave rise to the selection of broad-spectrum antibiotics in most of the cases. This is also due to the fact that overall susceptibility is modelled by means of a logical-AND gate. However, prescription of unnecessary broad-spectrum antibiotics is undesirable from a medical point of view, as it stimulates the development of antibiotic resistance in pathogens. We subsequently adopted the closed-world assumption in interpreting sputum-culture data when filling in evidence in the Bayesian network, which resulted in an improved therapeutic performance of the model. When using the CWA, the percentage of patients for whom a correct antibiotic therapy was advised by the Bayesian network, improved by 50%. We conclude that the CWA, which was originally developed in the context of relational database theory, may also be useful when interpreting clinical data using a Bayesian network.

In the near future, we plan to investigate ways to

ISSN: 1727-1983

Table 3: Absolute and relative numbers of incorrectly and correctly advised antibiotics for both the original and new Bayesian-network model interpretation.

Model	**Incorrect**		**Correct**
	too narrow	too broad	
Original	6 (4%)	129 (82%)	22 (14%)
New (CWA)	23 (14%)	31 (20%)	103 (66%)

model interaction among variables in a more natural way using causal independence models [6].

Acknowledgements

This research has been partially supported by the Netherlands Organization for Scientific Research (NWO: grant 634.000.026) under the TOKEN-2000 programme, as part of the TimeBayes project.

References

[1] M.J.M. Bonten. Prevention of infection in the intensive care unit. *Current Opinion in Critical Care*, 10(5), 2004.

[2] P.J.F. Lucas, N.C. de Bruijn, K. Schurink, and A. Hoepelman. A probabilistic and decision-theoretic approach to the management of infectious disease at the icu. *Artificial Intelligence in Medicine*, 19(3):251–279, 2000.

[3] M.R. Genesereth and N.J. Nilsson. *Logical Foundations of Artificial Intelligence*. Morgan Kaufmann, 1987.

[4] M.J.M. Bonten, M.H. Kollef, and J.B. Hall. Risk factors for ventilator-associated pneumonia: from epidemiology to patient management. *Clinical Infectious Diseases*, 38(8), 2004.

[5] H.C Sox, M.A. Blatt, M.C. Higgins, and K.I. Marton. *Medical Decision Making*. Butterworth-Heinemann, 1988.

[6] P.J.F. Lucas. Bayesian network modelling through qualitative patterns. *Artificial Intelligence*, 163:233–263, 2005.

[7] D. Heckerman and J.S. Breese. Causal independence for probabilistic assessment and inference using bayesian networks. *IEEE Transactions on Systems, Man and Cybernetics*, 26(6), 1996.

[8] R. Reiter. On closed world databases. In H. Gaillaire and J. Minker, editors, *Logic and Data Bases*, pages 55–76. New York: Plenum Press, 1978.

ISSN: 1727-1983