A Dynamic Bayesian Network for Diagnosing Ventilator-Associated Pneumonia in ICU Patients

Theodore Charitos, Linda C. van der Gaag,
Inst. of Inform. and Comp. Sciences
Utrecht University
P.O Box 80.089
3508 TB Utrecht, The Netherlands
theodore, linda@cs.uu.nl

Stefan Visscher, Karin Schurink,
Dept. of Internal Medic. and Infect. Diseases
University Medical Centre Utrecht
Heidelberglaan 100
3584 CX Utrecht, The Netherlands
S.Visscher, K.Schurink@azu.nl

Peter Lucas
Inst. for Comp. and Inform. Sciences
Radboud University, Nijmegen
Toernooiveld 1
6525 ED Nijmegen, The Netherlands
peterl@cs.ru.nl

Abstract

Diagnosing ventilator-associated pneumonia in mechanically ventilated patients in intensive care units is currently seen as a clinical challenge. The difficulty in diagnosing ventilator-associated pneumonia stems from the lack of a simple yet accurate diagnostic test. To assist clinicians in diagnosing and treating patients with pneumonia, a decision-theoretic network was designed with the help of domain experts. A major limitation of this network is its inability to represent pneumonia as a dynamic process that progresses over time. In this paper, we construct a dynamic Bayesian network that explicitly captures the development of the disease through time. We discuss how probability elicitation from domain experts serves to quantify the dynamics involved and show how the nature of patient data helps reduce the computational burden of inference. We evaluate the diagnostic performance of our dynamic model and report promising results.

1 Introduction

Many patients admitted to an intensive care unit (ICU) need respiratory support by a mechanical ventilator; in addition, many of these patients are affected by severe disease which may result in depression of their immune system. Both conditions promote the development of ventilator-associated pneumonia (VAP) in these patients. Because of the wide-spread dissemination of multiresistant bacteria at the ICU, effective and fast treatment of VAP is seen as an issue of major significance. The difficulty of the diagnosis of VAP is in the lack of a gold standard; VAP is therefore diagnosed by taking a number of different clinical features into account [7].

A probabilistic and decision-theoretic network [3], representing the uncertainties and preferences involved in dealing with the treatment of VAP, was constructed by Lucas et al. [4]. The network was developed with the help of two infectious disease experts, who assessed both its qualitative structure and its numerical part. The goal of the network was to prescribe an optimal antimicrobial therapy, and thereby assist clinicians in treating patients with VAP.

A prominent role in the domain of pneumonia is played by two stochastic processes: the *colonisation* of the laryngotracheobronchial tree by pathogens and the onset and development of *pneumonia*. Although both processes evolve dynamically, these dynamics were not explicitly modelled by means of temporal transitions in the network described above. Instead, the dynamics of the processes were implicitly modelled by additional interactions between the duration of stay and the duration of mechanical ventilation of a patient with the colonisation by pathogens. The main motivation for this simplification was the large amount of data needed to specify the probability distribution underlying the stochastic processes and the increase in computational requirements. The network thus constitutes a *static* simplification of the domain which obscures its dynamic nature. In fact, the static network was used for every patient for each day on the ICU separately, without taking into account the patient's characteristics from earlier days. As the development of VAP is a dynamic process, we need to model time in a more explicit way to improve the diagnosis.

In this paper, we ameliorate the problems related with having modelled VAP as a dynamic process. We develop a dynamic Bayesian network that explicitly captures the temporal relationships between the variables [5]; our focus thereby is on the diagnostic part of the network. We use the method of Van der Gaag et al. [9], for the elicitation, from domain experts, of the probability distribution of the underlying stochastic process. This method transcribes probabilities and uses a scale with both numerical and verbal anchors that assists experts to assess many probabilities in little time. Moreover, we discuss how the computational burden of inference in our model can be eased by exploiting the nature of the observations involved, with just a small loss in accuracy [2].

We evaluated our dynamic network on a group of patients, drawn from the files of the ICU of the University Medical Centre Utrecht in the Netherlands. Our results indicate that the dynamic model is capable of distinguishing between patients with VAP and without VAP. By exploiting all available past information of a patient, it in fact yields better predictions than the static model. This occurs specifically for patients without VAP, for whom we notice that the use of previous information leads to much lower estimates for VAP than the ones obtained from the static network.

The remainder of this paper is organised as follows. In the next section, we briefly describe the static probabilistic and decision-theoretic network that had been developed be-

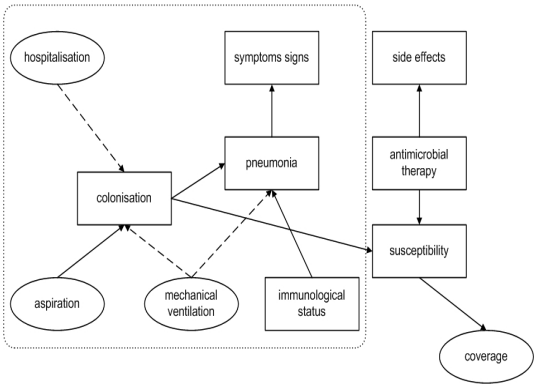

Figure 1: Global structure of the sVAP network. The dashed box indicates the network's diagnostic part.

fore for the management of patients with VAP. In Section 3, we discuss the construction of a dynamic network for VAP. Section 4 presents the results of an experimental evaluation of our network. The paper ends with our conclusions and directions for further research in Section 5.

2 Pathophysiology of VAP

Ventilator-associated pneumonia is a low-prevalence disease occurring in mechanically-ventilated patients in critical care and involves infection of the lower respiratory tract [1]. In contrast to infections of more frequently involved organs (such as the urinary tract), for which mortality is low, ranging from 1 to 4%, the mortality rate for VAP ranges from 24 to 50% and can reach 76% for some high-risk pathogens. Important factors related to the development of VAP include an increased *body temperature*, the use of *antipyretic drugs*, an abnormal amount of coloured *sputum*, *signs* on the chest X-ray, an abnormal ratio between the amount of oxygen in the arterial blood and the fractional inspired oxygen concentration, that is, pO_2/FiO_2, the duration of *mechanical ventilation*, and an abnormal number of *leucocytes*. As diagnosing a disorder in medicine involves reasoning with uncertainty, a decision-theoretic network was constructed as part of a decision-support system to assist clinicians in the diagnosis and treatment of VAP in the ICU [4],[7]. Figure 1 illustrates the network, which we call the static VAP network, os sVAP network for short. The signs and symptoms included in the sVAP network are shown in more detail in Figure 2.

The relationship between the *colonisation* by pathogens and the development of *pneumonia* is captured in the sVAP network as follows. Periodically, a sample of the patient's sputum is cultured at the laboratory. When the culture shows a number of colonies of a particular bacterium that is above a certain threshold, the patient is said to be colonised by this bacterium. The seven groups of microorganisms that occur most frequently in critically ill patients and cause colonisation, are modelled in the therapeutic part of the network. Figure 3 depicts the probabilistic relation between the seven groups of microorganisms of colonisation to pneumonia. Information about which bacterium or bacteria are currently present in a patient and the current signs

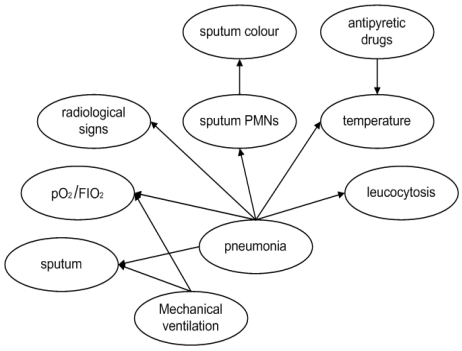

Figure 2: Symptoms and signs of pneumonia.

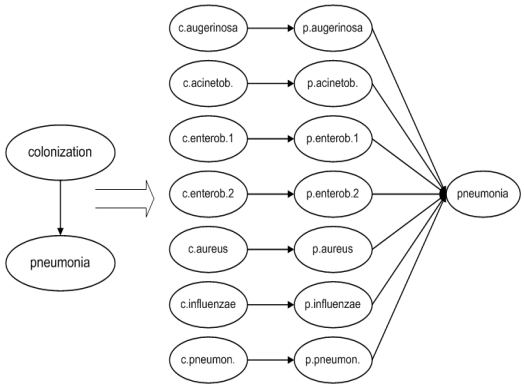

Figure 3: Detailed structure of the influence of colonisation on pneumonia.

and symptoms constitute the basis for choosing optimal antimicrobial treatment on multi-resistant bacteria and is considered best practice.

3 A dynamic Bayesian network for VAP

In this section, we describe the construction of a dynamic Bayesian network that represents explicitly the development of pneumonia. In addition, we address the computational burden of inference with the network.

3.1 Preliminaries

A dynamic Bayesian network is a graphical model that encodes a joint probability distribution on a set of stochastic variables, explicitly capturing the temporal relationships between them. More formally, let $\mathcal{V}_n = (V_n^1, \ldots, V_n^m), m \geq 1$, denote the set of variables at time step n. Then, a dynamic Bayesian network is a tuple (B_1, B_2), where B_1 is a Bayesian network that represents the prior distribution for the variables at the first time step $\mathcal{V}_1$, and B_2 defines the transitional relationships between the variables in two consecutive time steps, so that for every $n \geq 2$

$$p(\mathcal{V}_n \mid \mathcal{V}_{n-1}) = \prod_{i=1}^{m} p(V_n^i \mid \pi(V_n^i))$$

where $\pi(V_n^i)$ denotes the set of parents of V_n^i, for $i = 1, \ldots, m$.

We distinguish between two types of relationship in a dynamic Bayesian network: *transitional* relations that capture a dependence among variables between different time steps, and *local* relations that capture a dependence between variables within the same time step. If a relationship exists between the same variable over different time steps, this variable is called *persistent*. Based on the two types of relationship, per time step, the set of variables $\mathcal{V}_n$ is split into three mutually exclusive and collectively exhaustive sets $\mathcal{I}_n, \mathcal{X}_n, \mathcal{Y}_n$, where the sets $\mathcal{I}_n, \mathcal{Y}_n$ constitute the input and output variables and $\mathcal{X}_n$ consists of the hidden variables for the time step under study. Usually, $\mathcal{I}_n$ includes observable variables that affect the probability distribution of $\mathcal{X}_n$, while $\mathcal{Y}_n$ includes observable variables whose probability distribution is affected by $\mathcal{X}_n$. The set $\mathcal{X}_n$ includes the variables that represent the stochastic processes of the network and whose values are never observed. Later in the paper, we will need the notion of *forward interface* of a dynamic network, which is the set of variables at time step n that affect some variables at time step $n+1$.

Dynamic Bayesian networks are usually assumed to be time invariant, which means that the topology and the parameters of the network per time step and across time steps do not change. Moreover, the Markov property for transitional dependence is assumed, which means that $\pi(V_n^i)$ can include variables either from the same time step n or from the previous step $n-1$, but not from earlier time steps [5]. Then, by unrolling B_2 for N time steps, a joint probability distribution $p(\mathcal{V}_1, \ldots, \mathcal{V}_N)$ is defined for which the following decomposition property holds:

$$p(\mathcal{V}_1, \ldots, \mathcal{V}_N) = \prod_{n=1}^{N} \prod_{i=1}^{m} p(V_n^i \mid \pi(V_n^i))$$

Applying a dynamic Bayesian network usually amounts to computing the marginal probability distributions of the hidden variables at different times. The computations involved constitute the *inference*. Three types of inference are distinguished. *Monitoring* is the task of computing the probability distribution for $\mathcal{X}_n$ at time n given the observations that are available up to and including time n. *Smoothing* is the task of computing the marginal probability distribution for $\mathcal{X}_n$ at time n given the observations available up to time N where $N > n$. Finally, *forecasting* is the task of predicting the probability distribution of $\mathcal{X}_n$ at time n given the observations that are available about the past up to time N where $N < n$.

3.2 Modelling issues

A natural extension of the diagnostic part of the sVAP network is a network that represents time explicitly [4]. Figure 4 gives an overview of the structure of the dynamic network that we constructed for the diagnosis of VAP, which we call the dVAP network. The dVAP network includes two interacting dynamic hidden processes, modelled by the variables *colonisation* and *pneumonia*; there is no transitional influence between them, but both are persistent. The process of colonisation is influenced by three input variables, *hospitalisation*, *mechanical ventilation* and *previous antibiotics*, which in essence control its dynamics. We note

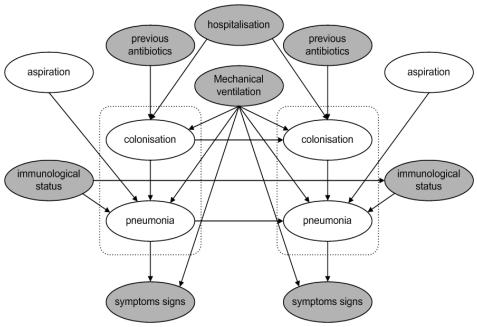

Figure 4: The dVAP network for the diagnosis of VAP; clear nodes are hidden, shaded nodes are observable. The dashed boxes indicate the hidden processes of the network.

that the variables *hospitalisation* and *mechanical ventilation* are observed for a period that is longer than the transition interval of the model. The variables thus are modelled as affecting adjacent time steps. The variable *previous antibiotics* represents the effect of previous medication to the patient on the process of colonisation.

One of the difficulties in constructing the dVAP model, was defining the length of the transition interval. It may seem trivial in general to decide upon the actual interval length, but in our case it proved to be rather difficult since there was no *a-priori* commonly acknowledged interval length that appropriately represents the evolution of the unobserved disease. Also, there was not a standard interval with which observations were collected in our data files. The latter can be attributed to most of the measurements being collected by nurses; for example, observable variables such as *body temperature* and *sputum colour* were measured frequently (approximately every two or three hours), while variables such as *radiological signs* and *leucocytosis* were measured once per day. Based on these insights, we decided to use a transition interval of one day (24 hours) for the dVAP network. Within this interval, the network *aggregates* the observations in a way similar to the previously constructed static network. For each observable variable, the value most frequently observed during the day was chosen as representative for that day; in cases where there was no prevalent value in the data, the worst value observed for the patient was chosen, to allow for *conservative* conclusions from the network. The chosen transition interval appealed to be compatible with the application characteristics and admissible by the domain experts.

A main issue in building the dVAP network was the acquisition of all conditional probabilities required. Although the three ICUs that acted as a setting for this study used the same shared computer-based patient record system, it appeared very hard to select relevant patient cases from the collected data. The main reason was that VAP is always a concomitant disease. As a consequence, clinicians tend to not report the presence of VAP in a patient. We thus found that only in a very small proportion of cases, a patient was reported as having VAP. Since we could not exploit the data for estimating the probabilities for our network, all parameters had to be assessed by experts. Compared to the sVAP network, the new parameters to be assessed concerned the dynamics of the stochastic processes of colonisation and

Suppose a patient has been mechanically ventilated for 48 hours and now has pneumonia caused by *s.aureus*. If this patient after 24 hours is *not mechanically ventilated*, but is *colonized with s.aureus* and has *phagocyte dysfunction*, then how likely is it that the patient will still have pneumonia caused by s.aureus ?

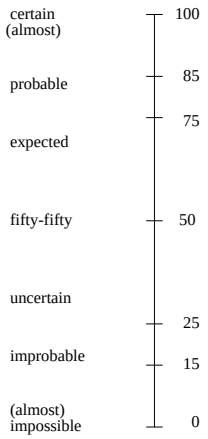

Figure 5: The fragment of text and probability scale for the assessment of the conditional probability $p(\textit{pneum.aureus=yes} \mid \textit{pneum.aureus=yes}, \textit{mech.ventilation=no,colonisation.aureus=yes, phagocytes.dysfunction=yes})$

pneumonia. To estimate those probabilities from domain experts we used the elicitation method proposed by Van der Gaag et al. [9]. This method is tailored to eliciting a large number of probabilities in a short time. Its main characteristic is the idea of presenting conditional probabilities as fragments of text and of providing a scale for marking assessments with both numerical and verbal anchors; for every conditional probability that needs to be assessed the domain experts are provided with a separate figure with the text and associated scale. Figure 5 shows, as an example, the figure pertaining to the conditional probability

$$p(\textit{pneum.aureus=yes} \mid \textit{pneum.aureus=yes}, \textit{mech.ventilation=no}, \textit{colonisation.aureus=yes}, \textit{phagocytes.dysfunction=yes})$$

for the dVAP network. In total, 2226 probabilities were elicited from a single domain expert within a few hours.

3.3 Computational issues

The practicability of the dVAP network depends to a large extent on the computational burden of inference with the network. For diagnosing patients with VAP, we monitor them at each time step. For this purpose, we use the *interface algorithm* with the dVAP network [5]. The interface algorithm is an extension of the *junction-tree algorithm* for inference with Bayesian networks in general [3], efficiently exploiting the forward interface of a dynamic network. The algorithm is linear in the total number of time steps and for large time scopes, the computation time can prove to be prohibitive for practical purposes.

Recent results show that, in case consecutive similar observations are obtained, the probability distribution of the hidden process converges to a limit distribution within a given level of accuracy [2]. After some number of time steps, therefore, there is no need for further inference as long as similar observations are obtained. The phenomenon of consecutive similar observations was evident for several patients in the ICU files. For example, for many patients we found that the same configuration of values was observed for all or almost all of the observable variables for a number of consecutive days.

Using the *relative entropy* distance measure for distributions, we can further show that it suffices to use just the most recent data for monitoring. This result depends on the properties of the transition matrix that models the evolution of the process, but a detailed description is out of the scope of the present paper. We define the *forward acceptable window* $\omega^f_{n,\epsilon}$for the present time step n given a specified level of accuracy ϵ, to be the minimal number of time steps that we need to use from the past to compute the probability distribution of the hidden variable at the present time within the level of accuracy ϵ. The scheme below illustrates the concept of the forward acceptable window, whose value can be established based upon the properties reviewed above:

$$\underbrace{\{1, \ldots, n_f, \ldots, n\}}_{\text{total time scope}} \longrightarrow \underbrace{\{n_f, \ldots, n\}}_{\omega^f_{n,\epsilon}}$$

We can now perform inference for time step n by considering only the forward acceptable window $\omega^f_{n,\epsilon}$ without losing too much in accuracy. Note that by doing so, the runtime requirements decrease from $O(n)$ to $O(n - n_f)$.

The main conclusion from the above considerations is that monitoring in the dVAP network can be eased considerably by exploiting the nature of the observations for a patient and by using the forward acceptable window.

4 Diagnostic performance

We evaluated the performance of the dVAP network, focusing on its diagnostic prediction per day. At our disposal we had a temporal database with data from 2233 distinct patients. Each record contains data collected for a patient during a one day stay in the ICU. The source of these data is the clinical management system used at the Intensive Care Units of the University Medical Centre Utrecht in the Netherlands. For 157 of these patients, VAP was established by two infectious-disease specialists. The conclusions obtained by the dVAP network were examined on a group of 20 patients in total, 5 of which were diagnosed with VAP. For these 5 patients we used the data from the day of admission to the ICU until the day they were diagnosed with VAP which was 10 days per patient. For each of these 5 patients, three patients for whom it was known

symptoms	VAP n = 5	no VAP n = 15
abnormal temperature	60%	7%
mech. ventilation (mean)	10d	10d
abnormal leucocytes	80%	53%
abnormal pO_2/FiO_2	60%	27%
abnormal sputum	80%	73%
coloured sputum	60%	60%
colonised	40%	13%
antipyretic drugs	100%	87%
positive X chest	40%	0%

Table 1: Data summary

that they did not develop VAP over time, were matched on three criteria: gender, number of mechanically ventilated days, and ICU ward. Table 1 summarises the data for the 5 patients with VAP and for the 15 patients without VAP at the tenth day of admission.

To compare the diagnostic performance of the dVAP network to that of the original sVAP network, we used the Brier score [6], [8]. We illustrate the Brier score for our dVAP network. For each patient i, the network yields a probability distribution p_i over the two values $j = 1, 2$ (yes, no) of pneumonia. The Brier score B_i for this distribution is defined as

$$B_i = \sum_{j=1,2} (p_{ij} - s_{ij})^2$$

where $s_{ij} = 1$ if the medical record of the patient states the value j, and $s_{ij} = 0$ otherwise. If the network would yield the correct value with certainty for a patient, then the associated Brier score would be equal to 0. For the probability distribution computed for any patient, therefore, the Brier score ranges between 0 and 2, and the better the prediction is, the lower the score. The Brier scores for all patients as well as the probability of VAP at the day it was diagnosed, for the dVAP and the sVAP networks respectively, are shown in Table 3. We note that for 15 patients of the total of 20 the computed Brier score was lower with the dVAP model than with the sVAP network.

The quality of the two networks can be expressed in an overall score that is computed from the scores for our collection of patients. For m patients, the overall Brier score is defined as

$$B = \frac{1}{m} \sum_{i=1,\ldots,m} B_i$$

The overall Brier score for the sVAP network can be readily computed from Table 3 and equals 0.3370, while the overall Brier score for the dVAP network is 0.2376. The lower score for the dVAP network conveys the information that this network is better informed than the sVAP network and can arrive at relatively good estimates for diagnosing VAP.

Compared to the sVAP network, the dVAP network takes into consideration the history of a patient. For the patients $22122, 23844, 24114, 21542, 22736$ for example, who were not diagnosed with VAP, the dVAP network derived low probabilities for the presence of VAP by exploiting all previous information. The sVAP network, in contrast, used just the current information and produced much higher probabilities. For the patients diagnosed with VAP, the two models behave more or less similarly, with the highest absolute discrepancy observed in patient 28393, to whom the sVAP network assigned a probability of VAP of 0.997 and the dVAP network assigned a probability of VAP of 0.904.

patient id.	**24528**	**22303**	**23505**	**23844**
exact	0.9987	0.0015	0.0005	0.0325
$\omega^f_{10,0.003}$	0.9987	0.0013	0.0005	0.0347

Table 2: Exact and approximate probabilities for VAP for a group of matched patients.

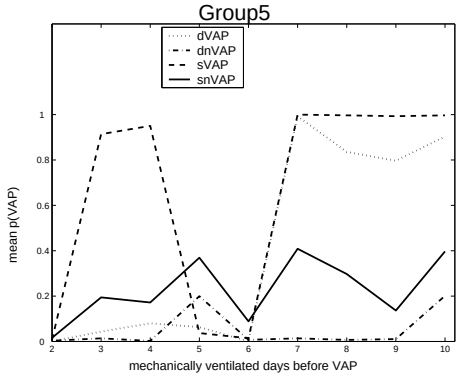

Figure 6: The dVAP and sVAP performance over time for a group of matched patients; dnVAP and snVAP represent the performance for the three patients without VAP combined.

To study the performance of the dVAP network over time, we computed the probability of VAP for each day and compared it to the respective probability from the sVAP network. In Figure 6 we plot, for a single group of four related patients, the probability of VAP for patient 28393 and the mean probability of VAP for the matched patients $21542, 22301, 22736$, from both networks. We observe that for the patient with VAP the trend in both networks is more or less the same after the fifth time step; to the patients without VAP, however, the dVAP network assigns lower probabilities than the sVAP network. The dVAP network thus is better able to distinguish between VAP and non-VAP patients.

To conclude, we performed the computations in the dVAP network using different values for the forward acceptable window $\omega^f_{n,\epsilon}$. We conclude that instead of using the observations for all 10 days in the ICU to compute the probability of VAP, we can use the observations for just the last 5 days with an average error for all patients smaller than $\epsilon = 0.003$. For a particular group of matched patients for example, the exact and approximate probabilities for VAP are showed in Table 2. We can thus use this forward acceptable window to speed up the computations and obtain results with an almost negligible error.

5 Discussion

In this paper, we discussed the construction of a probabilistic model that is aimed at assisting ICU clinicians in diag-

patient id.	VAP	sVAP	sBrier	dVAP	dBrier
22022	yes	0.996913	$1.90591 \cdot 10^{-5}$	0.9987	$3.38 \cdot 10^{-6}$
22563	no	0.0203017	$8.24318 \cdot 10^{-4}$	0.1395	0.0389205
22716	no	0.167208	0.055917031	0.0558	0.00622728
22730	no	0.00276365	$1.52755 \cdot 10^{-5}$	0.0002	$8 \cdot 10^{-8}$
23397	yes	0.00972048	1.961307055	0.0002	1.99920008
22122	no	0.430888	0.371328937	0.0316	0.00199712
22634	no	0.0203017	0.000824318	0.0003	$1.8 \cdot 10^{-7}$
22659	no	0.193411	0.07481563	0.0309	0.00190962
24528	yes	0.999959	$3.362 \cdot 10^{-9}$	0.9987	$3.38 \cdot 10^{-6}$
22303	no	0.0226662	0.001027513	0.0015	$4.5 \cdot 10^{-6}$
23505	no	0.0457446	0.004185137	0.0005	$5 \cdot 10^{-7}$
23844	no	0.297688	0.177236291	0.0325	0.0021125
25724	yes	0.0347989	1.863226327	0.0033	1.98682178
23872	no	0.0203017	0.000824318	0.0005	$5 \cdot 10^{-7}$
24114	no	0.43644	0.380959747	0.099	0.019602
24151	no	0.00999126	$2.22311 \cdot 10^{-5}$	$7 \cdot 10^{-8}$	$9.8 \cdot 10^{-15}$
28393	yes	0.996666	$2.22311 \cdot 10^{-5}$	0.9035	0.0186245
21542	no	0.175202	0.061391482	0.0218	0.00095048
22301	no	0.0740135	0.010955996	0.0013	$3.38 \cdot 10^{-6}$
22736	no	0.942073	1.775003075	0.581	0.675122

Table 3: Brier scores for the sVAP network and for the dVAP network, respectively.

nosing ventilator-associated pneumonia. In contrast to previous approaches that used a static decision-theoretic network for this low-prevalence disease, we focused on its dynamic evolution and used a dynamic Bayesian network as the primary tool for representation and inference.

We detailed various modelling steps in the construction of our dynamic network and described the use of an efficient procedure for expert elicitation of the probabilities required. We further argued that a number of convergence properties of dynamic Bayesian networks can be exploited to arrive at feasible algorithms that restrict the computational burden of inference with such a model. In this way, we ameliorated two important problems that were considered impervious in the past: the specification of the probabilities underlying the stochastic process modelled in the network and the computational burden of inference.

We evaluated our network on a set of ICU patients to examine its diagnostic accuracy. The lower overall Brier score of the dynamic network in comparison to the static one, indicated that representing time explicitly and taking into consideration the history of the patient, increases diagnostic performance. In our evaluation experiments, the dynamic network proved to be better at distinguishing between VAP and non-VAP patients than the static network, especially by assigning lower probabilities of VAP to the non-VAP patients. In the near future, we intend to improve the dVAP network by use of the available data for parameter learning and to test it on more ICU patients with the aim of embedding it in the clinical information system of the ICU.

Acknowledgements

This research was (partly) supported by the Netherlands Organization for Scientific Research (NWO). The authors would like to thank Marc Bonten for his valuable comments on an earlier version of this paper.

References

[1] M.J. Bonten. Prevention of infection in the intensive care unit. *Current Opinion in Critical Care*, 10(5):364-368, 2004.

[2] T. Charitos, P. de Waal, and L.C. van der Gaag. Speeding up inference in Markovian models. *Proceedings of the 18th International FLAIRS conference*, pp. 785-790, 2005.

[3] R.G. Cowell, A.P. Dawid, S.L. Lauritzen, and D.J. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Springer, 1999.

[4] P.J.F. Lucas, N.C de Bruijn, C.A.M 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.

[5] K. Murphy. *Dynamic Bayesian networks: Representation, Inference and Learning*. PhD thesis, University of California Berkley, 2002.

[6] H.A. Panofsky and G.W. Brier. *Some applications of Statistics to Meteorology*. The Pennsylvania State University, University Park, Pennsylvania, 1968.

[7] C.A.M. Schurink. *Ventilator Associated Pneumonia: a Diagnostic Challenge*. Ph.D thesis, Utrecht University, 2003.

[8] L.C. van der Gaag and S. Renooij. Probabilistic networks as probabilistic forecasters. *Proceedings of the Ninth Conference on Artificial Intelligence in Medicine in Europe* pp. 294-298, 2003.

[9] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B. Aleman, and B.G. Taal. How to elicit many probabilities. *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, pp. 647-654, 1999.