

Bayesian Networks for Cardiovascular Monitoring

Jennifer M. Roberts, Tushar A. Parlikar, Thomas Heldt, and George C. Verghese

Abstract—Bayesian Networks provide a flexible way of incorporating different types of information into a single probabilistic model. In a medical setting, one can use these networks to create a patient model that incorporates lab test results, clinician observations, vital signs, and other forms of patient data. In this paper, we explore a simple Bayesian Network model of the cardiovascular system and evaluate its ability to predict unobservable variables using both real and simulated patient data.

I. INTRODUCTION

Physicians have access to many types of information when treating patients. For example, they can examine real-time waveform data like blood pressure and ECG recordings; data trends like time-averaged heart rate; intermittent measurements like temperature and lab results; and qualitative observations like reported dizziness, nausea, or skin color. Within the Intensive Care Unit (ICU), physicians attempt to consider as much of the relevant information as possible, but the astronomically large amounts of data collected make it impossible to consider all available information within a reasonable amount of time. In addition, not all of the data collected is helpful in its raw form, but sufficient statistics taken from such data might help physicians gain a more thorough understanding of recent changes in the patient's state. Because of this, we are exploring ways to integrate different types of patient data into more synthesized forms (see <http://mimic.mit.edu/>).

A. Cardiovascular Models

Models provide one way of synthesizing multiple observations of the same complex system. Many traditional cardiovascular models consist of sets of differential equations derived from physiology using mechanical models of the heart and blood vessels. These models vary in both the levels of complexity and levels of temporal resolution.

Such models pose a number of problems when applied in an ICU setting. Only a limited number of measured signals are consistently available in the ICU, so one cannot accurately estimate more than a small subset of physiological

model parameters using the available data [2]. In addition, quantitative models of this sort fail to use all available information, as they usually do not incorporate diagnostic information, information about medications, qualitative observations, intermittent data, and observations from the clinical staff.

In order to integrate both qualitative and quantitative patient information into a single model, we are exploring ways to synthesize Bayesian Network models with traditional physiological ones. Bayesian Networks can incorporate non-numeric, discrete, and continuous information into the same model, while providing a stochastic framework for estimating unobservable variables, even when incoming patient data becomes unreliable.

B. Bayesian Networks in Medicine

Within the medical field, Bayesian Networks have been used for diagnosis, prognosis, and treatment selection [3], [4], [6]. For instance, the Heart Disease Program [4] uses a several-hundred-node Bayesian Network to diagnose patients. While useful, this extremely large network is somewhat difficult to understand and maintain. Because of this, we are working to develop simpler models whose parameters can be learned and updated from patient data.

Berzuni et al. [5] proposed a methodology for using Bayesian Networks to monitor patients and aid in drug therapy. Their system incorporated both general population data and incoming patient data to provide patient-specific models.

VentPlan [6], a Ventilator Therapy Planner, uses a relatively simple Bayesian Network and a quantitative physiological model of the pulmonary system to monitor patients and suggest ventilation settings. We aim to develop a similar Bayesian Network model for the cardiovascular system and use it to estimate unavailable information about internal patient state.

II. BAYESIAN NETWORKS

Bayesian Networks represent joint probability distributions using directed acyclic graphs [7]. The network nodes represent random variables, while the network arrows, which point from parent nodes to child nodes, indicate conditional dependencies between nodes. A given graph structure represents a class of probability distributions that factors in a particular way and thus shares a set of conditional independencies. Within a Bayesian Network, the joint node distribution equals the product of the conditional distributions for each node, given the node's parents.

This work was supported in part by the National Aeronautics and Space Administration (NASA) through the NASA Cooperative Agreement NC 9-58 with the National Space Biomedical Research Institute and by grant RO1 EB001659 from the National Institute of Biomedical Imaging and Bioengineering of the National Institutes of Health. J.M. Roberts gratefully acknowledges financial support from the Fannie and John Hertz Foundation and the National Science Foundation.

The authors are with the Laboratory for Electromagnetic and Electronic Systems within the Electrical Engineering and Computer Science Department at the Massachusetts Institute of Technology, Cambridge, MA 02139, email: verghese@mit.edu




Fig. 1. Simple Bayesian Network model of the cardiovascular system.

The network in Fig. 1 describes a probabilistic relationship between the random variables stroke volume (SV), mean arterial blood pressure (BP), total peripheral resistance (TPR), cardiac output (CO), and heart rate (HR) such that

$$P(SV, BP, TPR, CO, HR) = P(TPR)P(HR)P(SV)P(BP|TPR, CO)P(CO|HR, SV). \quad (1)$$

This model reflects the relationships between cardiac output and the other random variables, namely that $CO = HR \times SV$ and $BP = TPR \times CO$.

Bayesian Networks imply a set of conditional dependencies and independencies based on the fact that each node is conditionally independent of its non-descendants, given its parents. Within Fig. 1, TPR and CO are conditionally dependent, given BP, and BP is conditionally independent of both HR and SV, given CO. Intuitively, this means that if BP is fixed, knowing the value of TPR yields additional information about CO. Similarly, if CO is fixed, knowing BP yields no additional information about HR or SV.

Efficient algorithms exist for calculating the *a posteriori* distribution of unknown variables on a Bayesian Network, given available information, so we can use these networks to obtain estimates of unknown variables, given known data.

In working with cardiovascular models, we aim to create simple models that capture necessary dynamics without adding extraneous model parameters that are difficult to estimate, and we explore how accurately such models capture patient dynamics. Each random variable in Fig. 1 is modeled using a discrete probability mass function (PMF) which assigns probabilities to five levels for each random variable values. PMF parameters are modeled using Dirichlet priors [7]. For the preliminary work presented in this paper, we tested the network's ability to estimate patient parameters, first when presented with simulated patient data and then when presented with real patient data.




Fig. 2. The simple pulsatile cardiovascular model uses a 3-way switch which allows for simpler analysis of the circuit. The model has a single ventricular compartment (R_1 , $C_h(t)$, R_2), an arterial compartment (C_a , R_3), a venous compartment (C_v), an interstitial fluid compartment (C_i , R_i), and a blood infusion or hemorrhage source Q_v .

III. SIMULATED PATIENT DATA

A. Methods

1) *Data Generation*: In this section, the data used to train and test the Bayesian Network was obtained by simulation of the simple pulsatile cardiovascular model introduced as in [1]. Fig. 2 illustrates the circuit representation for the model. C_a is the arterial compliance, C_v is the venous compliance, $C_h(t)$ is the time-varying compliance of a single ventricular chamber, R_1 is the inflow resistance to the ventricle, R_2 is the outflow resistance from the ventricle, and R_3 is the total peripheral resistance. The pressure V_h is the ventricular pressure, V_v is the central venous pressure, and V_a is the arterial blood pressure. The ventricular volume is Q_h . In addition, the model contains an interstitial fluid compartment with compliance C_i and resistance R_i , and a blood infusion or hemorrhage source Q_v .

The elastance function $E_h(t) = 1/C_h(t)$ for the ventricle is represented as a piecewise linear function [1] given by:

$$E_h(t) = \begin{cases} \frac{3(E_s - E_d)}{T}t + E_d & \text{for } 0 \leq t \leq \frac{T}{3} \\ \frac{6(E_s - E_d)}{T}(\frac{T}{3} - t) + E_s & \text{for } \frac{T}{3} \leq t \leq \frac{2T}{3} \\ E_d & \text{for } \frac{2T}{3} \leq t \leq T \end{cases} \quad (2)$$

where T is the duration of the cardiac cycle, E_s is the end-systolic elastance, and $E_d (\ll E_s)$ is the end-diastolic elastance. State equations for this model can be derived as in [1], but with an additional equation representing the interstitial compartment dynamics. The data was created using the same nominal parameters as in [1], combined with additional nominal parameters of $C_i = 300 \frac{ml}{mmHg}$, $R_i = 1 \frac{mmHg}{ml/sec}$, and $Q_v = 0 \frac{ml}{sec}$.

To generate testing and training data, we used the model to run twenty 500-second-long simulations during which one or more patient parameters deviated from their nominal values. These deviations are described in Table I. The cardiovascular responses, sampled at a rate of 10 Hz, were used to compute the beat-to-beat averaged values of the arterial and venous

TABLE I
DESCRIPTION OF THE TRAINING DATA SEGMENTS.

Data Segment	Description of parameter changes
1	Nominal Parameters
2	TPR increases from 1 to 5 PRU = mmHg/(ml/s)
3	$Q_v=4$ ml/s
4	$Q_v=-4$ ml/s
5	TPR decreases from 1 to 0.3 PRU
6	C_a decreases from 2 to 0.6 ml/mmHg
7	HR is stepped up from 60 to 300 bpm
8	HR is stepped down from 60 to 30 bpm
9	C_v decreases from 50 to 20 ml/mmHg
10	R_2 increases from 0.01 to 0.5 PRU
11	R_2 decreases from 0.01 to 0.001 PRU
12	E_s increases from 2.5 to 5 mmHg/ml
13	E_d decreases from 0.1 to 0.05 mmHg/ml
14	E_s decreases from 2.5 to 1 mmHg/ml
15	E_d increases from 0.1 to 0.2 mmHg/ml
16	$Q_v=4$ ml/s and TPR increases from 1 to 3 PRU
17	$Q_v=-4$ ml/s and TPR decreases from 1 to 0.5 PRU
18	HR is stepped up from 60 to 300 bpm TPR decreases from 1 to 0.5 PRU
19	HR is stepped down from 60 to 30 bpm TPR decreases from 1 to 3 PRU
20	HR is stepped up from 60 to 300 bpm TPR increases from 1 to 3 PRU

blood pressures, as well as the average flow across the total peripheral resistance. These variables were then used to calculate cardiac output. Using heart rate ($HR=\frac{60}{T}$) and CO, we calculated SV. These five variables, HR, SV, CO, BP, and TPR, were saved for each of the twenty simulations. The first ten samples of each simulation were discarded to remove unrealistic transient simulation effects, and the resulting data was then used to train and test the Bayesian Network.

2) *Training and Test Sets:* The data was first segmented into ten second intervals. Half of the intervals were chosen uniformly at random to be included in the training set, while the remaining intervals were placed in the test set. This random selection process was repeated fifty times to obtain fifty different training and test sets. This procedure was repeated for interval sizes of 100 seconds and 250 seconds.

Data was then quantized by placing values into five equally sized bins that cover an acceptable dynamic range for each variable of interest. HR was quantized to values of 25, 75, 125, 175, and 225 beats per minute (bpm).

3) *Network Training:* The network parameters are initialized using likelihood-equivalent uniform Dirichlet distributions [7]. The parameters of the Dirichlet priors are set to values less than one so that the resulting uniform prior is quickly overridden by incoming training data. The network training algorithm essentially computes histograms of the number of times that certain network variables take a given set of values simultaneously. These histogram counts are added to the parameters of the Dirichlet priors so that value combinations not encountered in the training data set are still assigned a small non-zero probability. The histogram counts summed with the prior parameters are normalized to create the new conditional probability distributions for each node. The algorithms can be found at <http://bnt.sourceforge.net/>.

We first trained a network on the entire data set to obtain a baseline for comparison. Estimates obtained from this

network reveal the optimal estimates for our data set. We then trained additional networks on each of the randomly segmented training sets described above.

4) *Testing the Trained Network:* We provided each trained network with quantized BP and HR values from the test set and used the network to obtain estimates of CO, TPR, and SV, given these values. We computed both quantized minimum mean square error (MMSE) estimates and maximum *a posteriori* (MAP) estimates as follows:

$$\begin{aligned} MMSE(X|BP, HR) &= \{E[X_q|BP_q, HR_q]\}_q \\ &= \{\sum x_q P(x_q|BP_q, HR_q)\}_q, \end{aligned} \quad (3)$$

and

$$MAP(X|BP, HR) = \arg \max_{x_q} \{P(x_q|BP_q, HR_q)\}, \quad (4)$$

where X is either CO, TPR, or SV; the subscript q indicates that the corresponding value is quantized; and x_q ranges over the set of quantized values taken by the random variable X. We then mapped the quantized estimates to their corresponding bins and compared the binned estimates with the actual binned values from the corresponding test set.

5) *Error Analysis:* We examined two types of error rates, the absolute error rate, i.e., the number of errors divided by the total number of samples, and the root mean square error. Any time an estimate did not fall in the same bin as the actual value, the estimate was counted as an error. The bin numbers were used to calculate the root mean square error, so that the error rate for each variable was normalized to the same scale. Bins were used to calculate the errors since identifying the relative range of the variable of interest is normally acceptable in a clinical setting.

B. Results

Error rates and standard deviations increased as the length of the training segment increased, since data combinations not seen in the training sets were more likely to appear in the test sets when the segment size was large. Analysis is presented for the 250-second segment case since those error rates are the most conservative.

Error rates for both types of estimators were comparable. As seen in Table II, the MAP estimator provided a slightly smaller absolute error rate, and the MMSE estimator provided a slightly smaller root mean square error, but these differences were not significant. If left to random chance, the expected absolute error rate would be approximately eighty percent, so the Bayesian Network is performing far better than random chance. The twenty-five to thirty-three percent absolute error rates indicate that the estimate falls in the correct bin more than two-thirds of the time. The small root mean square errors indicate that when the estimate does not fall in the same bin as the actual value, it tends to fall in a nearby or adjacent bin.

Errors more frequently occur when a combination of variable values appears in the test set but not the training set. Fig. 3 shows how such occurrences affect the learned conditional PMFs.

TABLE II
A COMPARISON OF THE MEAN ERROR RATES AND STANDARD DEVIATIONS FOR THE DATA SETS WITH 250 SECOND SEGMENTS

Estimator	Mean Absolute Error Rates (STD)		Mean Root Mean Square Error (STD)	
	MMSE	MAP	MMSE	MAP
Resistance	0.33 (0.06)	0.31 (0.06)	0.73 (0.14)	0.77 (0.17)
Cardiac Output	0.30 (0.07)	0.27 (0.07)	0.72 (0.16)	0.78 (0.23)
Stroke Volume	0.26 (0.08)	0.29 (0.06)	0.68 (0.16)	0.71 (0.14)

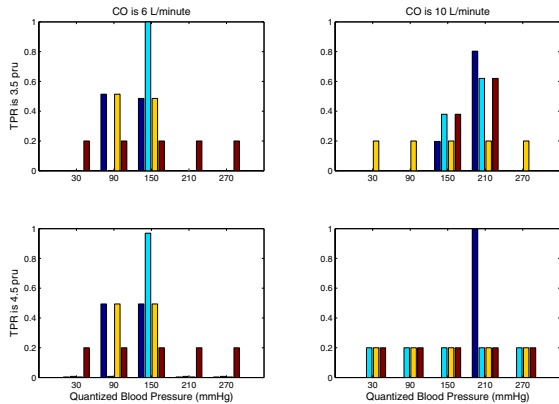


Fig. 3. Conditional probability mass functions (PMFs) for BP, given CO and TPR. Each plot shows the conditional PMF of BP for a particular set of CO and TPR values. Each set of bars of a given shading displays the PMF obtained when the network is trained using a different training set. The left-most bar in each group is obtained by training on the entire data set, while the other bars are obtained using partial training sets. When a particular BP-CO-TPR value combination is not found in a training set, the PMF remains uniform. The plot on the bottom right displays such a case.

IV. REAL PATIENT DATA

A. Methods

1) *Data*: As a preliminary test of the network’s ability to track real patient data, we presented the network with data from a single ICU patient record from the MIMIC II database [8]. Beat-to-beat estimates of cardiac output were calculated from the patient data using Liljestrand’s method [9] with the actual thermodilution cardiac output measurements for calibration. Beat-to-beat averaged BP waveforms were obtained from the 125-Hz blood pressure waveforms. We used the beat-to-beat BP, HR, and CO values to calculate beat-to-beat values for SV and TPR using $SV = \frac{CO}{HR}$ and $TPR = \frac{BP}{CO}$. All of the waveforms were then median filtered to reduce noise. The resulting waveforms were taken as the actual values of HR, BP, CO, TPR, and SV when training and testing the Bayesian network. Before testing and training, the data was quantized as in Sec. III-A.2. The resulting waveforms are pictured in Fig. 4.




Fig. 4. Patient data used to train the network. Original waveforms are shown in blue, quantized waveforms are shown in green, and the quantization thresholds are shown in grey.

2) *Network Training*: Training on one data set and testing on another resulted in unsatisfactory performance, so we used a sequential training method that allowed the network to continually update its probability distributions based on incoming patient data. We initialized the Dirichlet distributions as in Sec. III-A.3. The network then used sliding data windows of various sizes to set its probability distributions to normalized histograms of data from the most recent relevant 5000 beats plus the initialized Dirichlet parameters. The probability distributions for HR, SV, and TPR, Bayesian Network variables with no parents, depended on the most recent 5000 heart beats, approximately fifty-five minutes of data. Distributions for CO and BP, nodes with two parents, depended on the last 5000 heart beats for which the parent random variables took a particular set of values. For example, $P(CO|HR_b = 1, SV_b = 1)$, where HR_b and SV_b represent the bin numbers for HR and SV, respectively!, was computed from the last thousand beats for which $HR_b = 1$ and $SV_b = 1$. Data sets of this sort frequently extended further back in time than the most recent 5000 heart beats.

3) *Testing the Network*: During testing, the sequentially trained network uses the probability distribution learned from beats 0 through t to calculate its $(t + 1)^{st}$ estimates. After updating its probability distributions based on information from beat t , the network receives HR and BP information from beat $t + 1$ and uses this information to estimate the CO, SV, and TPR associated with the $(t + 1)^{st}$ beat.

4) *Error Analysis*: We compare the binned single beat-predictions with the binned actual values as in Sec. III-A.5.

B. Results

The MMSE estimates for TPR, CO, and SV are shown in Figure 5. Since the MMSE and MAP results were similar,




Fig. 5. Estimates when the Bayesian Network is sequentially trained with a training memory size of 5000 beats and tested on patient data. The original waveform is in blue, the quantized waveform in green, and the MMSE estimates in red.

we focused on the MMSE estimates. In general, the estimate tracked the original waveforms. When the data remained within the same set of quantization bins for long enough, however, the estimates settled to the corresponding quantized values, because the probability distribution windows no longer retained data points that took other values.

The sudden increase in SV at minute 1200 shows how the network responds to changes in the estimated variables that are not accompanied by changes in HR or BP. Without some type of indication from either HR or BP, the network cannot predict a sudden change in SV. Thus, the network adapts to the change in SV in a delayed fashion as the network updates its probability distributions based upon the new SV data. Because of the change in SV after minute 1200, the network successfully tracks a slight increase in the CO waveform, even though the quantized CO, HR, BP, and TPR values do not change. Here, the network relies on the changing probabilistic information about SV to change the CO estimate.

Within this data set, changes in SV and CO often occur without corresponding changes in BP and HR. Thus for this data set, the network frequently relies on information that the probability distributions extract from recent SV and CO combinations to estimate future CO and SV values.

The binned estimates for TPR equaled the actual binned value 99.7% of the time. The binned estimates for CO and SV were accurate 82% and 87.5% of the time respectively. Lower error rates were obtained by using smaller training window sizes, but the resulting estimates simply moved more quickly from one quantized value to another without improving the network's ability to approximate the actual waveform.

V. CONCLUSIONS AND FUTURE WORK

The network performance is promising. With proper training data, this model may provide satisfactory estimates of CO, SV, and TPR based on the BP and HR information generally available in the ICU. Future work will explore how to use both relevant population data and current patient data to evaluate the Bayesian Network probability distributions and to perform estimation. Procedures that incorporate general population data should be careful to rely only on population data relevant to the current patient, because using irrelevant data seems to hurt network performance. Future work should analyze how general population data might be used to improve network performance and under what conditions the use of general population data will yield unsatisfactory results. In addition, such work might explore how to incorporate additional patient information like lab results and intermittently available data into the Bayesian Network to aid the estimation process.

We also plan to focus on improving the network's ability to predict abrupt changes in the patient data. One can easily use recent data to approximate incoming patient data when the patient variables remain constant, but when a patient variable changes abruptly, we would like to use the Bayesian Network to combine patient history and relevant general population data to predict the patient's response to the sudden change. We are also exploring whether dynamic Bayesian Networks [7] might help capture such time correlations.

REFERENCES

- [1] T. Parlikar and G. Verghese, "A simple cycle-averaged model of cardiovascular dynamics," in *Proceedings of the 35th Annual Conference of the IEEE Engineering in Medicine and Biology Society*, September 2005, pp. 5490–5494.
- [2] Z. Samar, T. Heldt, G. Verghese, and R. Mark, "Model-based cardiovascular parameter estimation in the intensive care unit," in *Computers in Cardiology*, September 2005, pp. 635–638.
- [3] P. J. F. Lucas, L. C. van der Gaag, and A. Abu-Hanna, "Bayesian networks in biomedicine and health-care," *Artificial Intelligence in Medicine*, vol. 30, pp. 201–214, March 2004.
- [4] W. Long, "Temporal reasoning for diagnosis in a causal probabilistic knowledge base," *Artificial Intelligence in Medicine*, vol. 8, pp. 193–215, 1996.
- [5] C. Berzuini, R. Bellazzi, S. Quaglini, and D. J. Spiegelhalter, "Bayesian networks for patient monitoring," *Artificial Intelligence in Medicine*, vol. 4, pp. 243–260, May 1992.
- [6] G. Rutledge, G. Thomsen, I. Beinlich, B. Farr, B. Sheiner, and L. Fagan, "Combining qualitative and quantitative computation in a ventilator therapy planner," in *Proceedings of the Thirteenth Annual Symposium on Computer Applications in Medical Care*. Washington, D.C.: The American Medical Informatics Association, 1989, pp. 315–319.
- [7] R. E. Neapolitan, *Learning Bayesian Networks*. Upper Saddle River, NJ: Pearson Prentice Hall, 2004.
- [8] M. Saeed, C. Lieu, G. Raber, and R. Mark, "MIMIC II: A massive temporal icu patient database to support research in intelligent patient monitoring," *Computers in Cardiology*, vol. 29, pp. 641–644, 2002.
- [9] G. Liljestrand and E. Zander, "Vergleichende bestimmungen des minutenvolumens des herzens beim menschen mittels der stickoxydulmethode und durch blutdruckmessung," *Z Ges Exp Med*, vol. 59, pp. 105–122, 1928.