
Probabilistic modeling of sensor artifacts in critical care

Norm Aleks and Stuart Russell

Computer Science Division, University of California, Berkeley, CA 94720

NORM|RUSSELL@CS.BERKELEY.EDU

Michael G. Madden

Dept. of Information Technology, National University of Ireland, Galway, Ireland

MICHAEL.MADDEN@NUIGALWAY.IE

Kristan Staudenmayer, Mitchell Cohen,

Diane Morabito, and Geoffrey Manley

Dept. of Neurological Surgery, University of California, San Francisco, CA 94143-0112

STAUDENMAYER|MCOHEN@SFGHSURG.UCSF.EDU

MORABITOD|MANLEYG@NEUROSURG.UCSF.EDU

Abstract

We describe an application of probabilistic modeling and inference technology to the problem of analyzing sensor data in the setting of an intensive care unit (ICU). In particular, we consider the arterial-line blood pressure sensor, which is subject to frequent data artifacts that cause false alarms in the ICU and make the raw data almost useless for automated decision making. The problem is complicated by the fact that the sensor data are acquired at fixed intervals whereas the events causing data artifacts may occur at any time and have durations that may be significantly shorter than the data collection interval. We show that careful modeling of the sensor, combined with a general technique for detecting sub-interval events and estimating their duration, enables effective detection of artifacts and accurate estimation of the underlying blood pressure values.

1. Introduction

The work reported in this paper falls under the general heading of *state estimation*, i.e., computing the posterior distribution $P(\mathbf{X}_t | \mathbf{e}_{1:t})$ for the state variables \mathbf{X} of a partially observable stochastic system, given a sequence of observations $\mathbf{e}_{1:t}$. The specific setting for our work at the Center for Biomedical Informatics in Critical Care (CBICC) is an intensive care unit (ICU) at San Francisco General Hospital (SFGH) specializing in traumatic brain

Presented at the ICML/UAI/COLT 2008 Workshop on Machine Learning for Health-Care Applications, Helsinki, Finland, 2008. Copyright 2008 by the author(s)/owner(s).




Figure 1. One day's worth of minute-by-minute monitoring data for an ICU patient.




Figure 2. Arterial-line blood pressure measurement.

injury, part of a major regional trauma center. In this setting, the state variables \mathbf{X}_t include aspects of patient state, while the evidence variables \mathbf{E}_t include up to 40 continuous streams of sensor data such as blood pressures (systolic/diastolic/mean, arterial and venous), oxygenation of blood, brain, and other tissues, intracranial pressure and temperature, inspired and expired oxygen and CO_2 , and a wide array of variables measured by the mechanical ventilator.

A short section of data from these sensors is shown in Figure 1. This section illustrates a number of artifacts, including, in the top traces, sharp upward and downward deviations in blood pressure due to external interventions in the arterial line; in the middle traces, ubiquitous drop-outs in the venous oxygen level (blue line); and in the lower traces, many jagged spikes in measured lung compliance due to coughing.

The artifacts cannot be modeled simply as “noise” in the sensor model; many are extended over time (some for as long as 45 minutes) and most exhibit complex patterns of their own. Simple techniques for “cleaning” such data, such as median filtering, fail. Instead, we follow the general approach suggested by Russell and Norvig (2003), which involves careful generative modeling of sensor state using dynamic Bayesian networks (DBNs).

This paper focuses on the arterial-line blood pressure sensor (Figure 2), a key element of the monitoring system. As we describe in detail in Section 2, this sensor is subject to multiple artifacts, including drifting calibration error, discrete shifts due to manual recalibrations, and ar-

tificially high values due to the line flushes or the drawing of blood samples from the arterial line. These artifacts not only complicate the state estimation and diagnosis task; they also cause a large number of false alarms in the ICU, which lead in turn to true alarms being ignored and alarms being turned off (Tsien & Fackler, 1997). By modeling the artifact-generating processes, we hope to be able to infer the true underlying blood pressure even when artifacts occur.

To this point, the task described would be an applied Bayesian modeling problem of medium difficulty. What makes it slightly unusual and perhaps of more general interest is the fact that our sensor data are recorded at fixed intervals (one minute averages in our ICU) whereas the events of interest—in this case, re-zeroings, line flushes, and blood draws—can occur at any time and have durations ranging from under 5 seconds to over 100 seconds.

Thus, the natural time step for modeling the sensor state transitions might be one second, whereas the measurement interval is much larger. This brings up the question of how a “slow” (one-minute) model might be constructed and how it relates to a “fast” (one-second) model. This is an instance of a very important issue studied in the dynamical systems and chemical kinetics literatures under the heading of *separation of time scales* (see, e.g., Rao & Arkin, 2003). Fortunately, in our case the problem has a simple, exact solution: Section 3 shows that a one-minute model can be derived efficiently, offline, from the more natural one-second model and gives exactly the same evidence likelihood. The more general problem of handling multiple time scales within DBNs, noted by Aliferis and Cooper (1996), remains open.

Section 4 describes the complete model for blood pressure estimation, including artifact models, and Section 5 then evaluates the model on real patient data. We show a number of examples of artifacts, their detection, and inference of the underlying state values. We analyze model performance over more than 300 hours of data from 7 patients, containing 228 artifacts. Our results show very high precision and recall rates for event detection; we are able to eliminate over 90% of false alarms for blood pressure while missing fewer than 1% of the true alarms.

Our work is not the first to consider the probabilistic analysis of intensive care data. Indeed, one of the most well-known of the early Bayes net applications was the ALARM model for patient monitoring under ventilation (Beinlich et al., 1989)—although this model had no temporal element. The work most closely related to ours is that of Williams, Quinn, and McIntosh (2005), who apply factorial switching Kalman filters—a particular class of DBNs—to artifact detection in neonatal ICU data. Their (one-second) model is roughly analogous to the models described by




Figure 3. 1-second (top) and 1-minute-average (bottom) data for systolic/mean/diastolic pressures. On the left, a blood draw and line flush in quick succession. On the right, a zeroing.

Russell and Norvig, using Boolean state variables to represent events that block normal sensor readings. Another important line of work is the MIMIC project, which, like ours, aims to apply model-based methods to the interpretation of ICU data (Heldt et al., 2006).

2. Blood Pressure Monitoring

Blood pressure provides useful information about much of physiology and is typically measured continuously in the ICU. The most common ICU blood pressure measurement device is the *arterial line*, illustrated in Figure 2; a catheter placed into one of the patient’s small arteries is connected to a pressure transducer whose output is displayed on a bedside monitor.

Because blood flow varies during the cardiac cycle, blood pressure is pulsatile. In medical records, including our data set, blood pressure measurements are summarized in two or three values: *systolic* blood pressure, which is the maximum reached during the cardiac cycle, *diastolic*, which is the corresponding minimum, and sometimes the *mean*. (A derivative measure, the *pulse pressure*, is equal to the difference between the systolic and diastolic pressures.)

We consider the three common artifact types illustrated in Figure 3(top): 1) in a *blood draw*, sensed pressure gradually climbs toward that of the pressure bag, then suddenly returns to the blood pressure when the stopcock is closed, seconds or minutes later; 2) in a *line flush*, the transducer is exposed to bag pressure for a few seconds; 3) during *zeroing*, the transducer is exposed to atmospheric pressure (defined as zero). We refer to blood draws and line flushes collectively as “bag events.”

Figure 3(top) shows the artifacts using data collected at

one-second intervals. In our clinical setting, however, data are collected as one-minute average values, as shown in Figure 3(bottom). Each one-minute interval may be partially or wholly corrupted by one or more artifacts, and the recorded value is a linear function of the true pressure, the artifactual pressure(s), and the fraction of the minute occupied by artifact. Using systolic pressure s as an example, for an artifact of length p (as a fraction of the averaging interval) and mean artifact pressure x , the apparent systolic pressure $s' = px + (1 - p)s$.

Our DBN model in Section 4 includes summary variables and equations relating the one-minute readings to the true underlying pressures, artifacts’ durations, bag and atmospheric pressure, etc.; it can therefore estimate the duration and other characteristics of artifacts that have corrupted the data. Patterns produced by artifacts in the one-minute data are highly varied, but it turns out (see Section 5) that the detailed modeling pays off in revealing the characteristic relationships that follow from the nature of the corrupting events.

3. Modeling Sub-Interval Events

The data we work with are generated by a combination of physiological processes that vary over timescales of several minutes and artifactual events lasting perhaps only a few seconds. Thus, a natural choice would be to choose a “fast” time step for the DBN model, e.g., 1 second. On this timescale, the sensor state variables indicate whether or not an artifactual event is currently in progress. The transition model for these variables indicates the probability that a new event begins and the probability that it continues if already in progress. These probabilities can be estimated simply by measuring event frequencies and durations. For

memoryless (geometric) duration distributions, such as we find in Section 5, the model requires very few parameters.

The main drawback of using a fast time step is computational: inference must be carried out over 60 time steps for every one measurement that arrives. Furthermore, much of this inference is pointless given the lack of evidence at all the intervening time steps.

We could instead build a model with a “slow” time step of one minute, so that evidence arrives at each time step. The problem here is to determine the structure and parameters of such a model. First, to explain the evidence, we will need a count variable for each event type saying how many seconds of that minute were occupied by such events. Now, it is easy to see that such variables must depend on their corresponding variables one minute earlier: for example, if the preceding minute was fully occupied by a blood draw event, then that event was in progress at the beginning of the current minute, so the current minute is likely to be at least partially occupied by the event. Moreover, if the event types are mutually exclusive, then each count variable depends on *all* the preceding variables. As each count variable has 61 values, this leads to *huge* conditional distributions that summarize all possible ways that the preceding 60 seconds could be divided among the various event types. How could we possibly estimate these?

The astute reader will have noticed two things: first, adding a state variable for “event in progress at the minute boundary” will simplify the model somewhat. Second, the CPTs for the slow model can be *derived* from the fast model and need not be estimated or guessed. This is the typical situation with separation of time scales: slow-time-scale models are computationally more tractable but can only be constructed by deriving them from a fast-time-scale model. We now explain briefly how this is done for our problem.

Consider a fast model as shown in Figure 4(a). Let the fast time step be ϵ and a measurement interval be $N\epsilon$ (where $N=60$ in our domain). $f_i = 1$ iff an event is occurring at time $i\epsilon$; $G_{Nj} \equiv \sum_{i=N(j-1)}^{Nj-1} f_i$ counts the number of fast time steps within the j th measurement interval during which an event is occurring. The j th observed measurement E_{Nj} is determined entirely by G_{Nj} ; therefore, it suffices to consider the joint distribution over G_0, G_N, \dots, G_{Nt} .

To obtain a model containing only variables at the slow intervals, we simply need to sum out the f_i variables other than the ones at interval boundaries. We can do this topologically by a series of arc reversal and node removal operations (Shachter, 1986); a simple proof by induction (omitted) shows that, regardless of the number of fast steps per slow step, we obtain the reduced structure in Figure 4(b).




Figure 4. (a) DBN model showing relationships among the fast event variables f_i , interval count variables G_{Nj} , and measurement variables E_{Nj} . (b) A reduced model that has the same distribution for G_0, G_N, \dots, G_{Nt} .

By construction, this model gives the same joint distribution for G_0, G_N, \dots, G_{Nt} . Importantly, neither f_{Nj} nor G_{Nj} depends on $G_{N(j-1)}$.¹

To complete the reduced model, we need the conditional distributions $P(G_{Nj}|f_{N(j-1)})$ and $P(f_{Nj}|f_{N(j-1)}G_{Nj})$. That is, how many “ones” do we expect in an interval, given the event status at the beginning of the interval, and what is the probability that an event is occurring at the beginning of the next interval, given also the number of ones in the current interval? The fast model can be described with two parameters: $p = P(f_i = 1|f_{i-1} = 1)$, and $q = P(f_i = 1|f_{i-1} = 0)$. Given those parameters, conditional probability tables for the slow model can be calculated offline using dynamic programming, as follows. A table is constructed for the variables f_i and C_i for i from 1 up to N , where C_i is the number of ones up to $i - 1$ and $C_0 = 0$. The recurrences for f_i and C_i are as follows:

$$\begin{aligned}
 P(C_i, f_i = 1|f_0) = & \\
 & p P(C_{i-1} = C_i - 1, f_{i-1} = 1|f_0) + \\
 & q P(C_{i-1} = C_i, f_{i-1} = 0|f_0) \quad (1)
 \end{aligned}$$

$$\begin{aligned}
 P(C_i, f_i = 0|f_0) = & \\
 & (1 - p) P(C_{i-1} = C_i - 1, f_{i-1} = 1|f_0) + \\
 & (1 - q) P(C_{i-1} = C_i, f_{i-1} = 0|f_0) \quad (2)
 \end{aligned}$$

Extracting the required conditional probabilities from the table is straightforward. The table is of size $O(N^2)$, so the total time to compute it is negligible for any plausible value of N . Now we have the following result:

¹Intuitively, the distribution over G_{Nj} for the N th interval is determined by the value of f at the beginning of the interval, independent of $G_{N(j-1)}$, whereas f_{Nj} depends on the count G_{Nj} for the preceding interval because, for example, a high count implies that an event is likely to be in progress at the end of the interval.

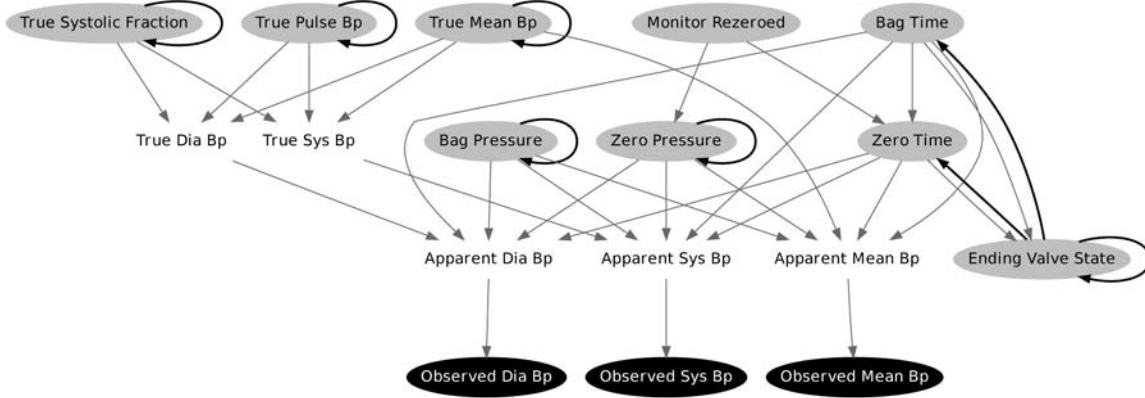


Figure 5. The blood pressure artifact detection DBN. Gray edges connect nodes within a time slice; black edges are between time slices. “Nodes” without surrounding ovals are deterministic functions included for clarity.

Theorem 1 *Given the conditional distributions computed by Equations 1 and 2, the reduced model in Figure 4(b) yields the same distribution for the count sequence G_0, G_N, \dots, G_{N_t} as the fine-grained model in Figure 4(a).*

The conditional distributions that we obtain by dynamic programming have some interesting limit cases. In particular, when events are short compared to measurement intervals and occur frequently, we expect the dependence on $f_{N(j-1)}$ to disappear and the distribution for G_{N_j} to be approximately Gaussian with mean $\frac{N}{1+p/(1-q)}$. When $p = q$, the f_i s become i.i.d. and G_{N_j} is exactly binomial—the recurrences compute the binomial coefficients via Pascal’s rule.

Generalizing the analysis to the case of multiple disjoint event types (i.e., f_i takes on more than two values) is mathematically straightforward and the details are omitted. There is, however, a complexity problem as the number of event types increases. The count variables G_{N_j} , H_{N_j} , and so on at time N_j are all dependent on each other given $f_{N(j-1)}$, and f_{N_j} depends on all of them; thus, using the approach given above, the precomputed tables will scale exponentially with the number of event types. This is not a problem in our application, where we do not expect sensors to have more than a few distinct types of “error” state. Furthermore, if each event type occurs independently of the others (except for the mutual exclusion constraint), then the conditional distribution for the count variable of each depends not on the *combination* of counts for the other types but on the *sum* of those counts, leading to low-order polynomial growth in the table sizes.

The preceding analysis covers only the case in which f_i depends just on f_{i-1} , leading to independently occurring events with a geometric length distribution. Constructing models with other length distributions is a well-studied

problem in statistics and most cases can be well approximated with a modest increase in the size of the dynamic programming table. Handling non-independent event occurrence is often more important; for example, blood draws may occur in clusters if multiple samples are required. Such dependencies can be handled by augmenting the state with timer variables, again at modest cost.

Before we move on to describe the complete model, it is important to note that a model with a finer time scale that the measurement frequency *can* provide useful extra information. By analogy with sub-pixel localization in computer vision, such a model can estimate the *time of occurrence* of an event within a measurement interval.

4. Combined model

The complete model for blood pressure measurements is shown in Figure 5. It has the same basic structure as the reduced model in Figure 4(b) but extends it in various ways.

The evidence variables \mathbf{E}_{N_j} are just the three reported blood pressure values *ObservedDiaBP*, *ObservedSysBP*, and *ObservedMeanBP*. These reflect, with some Gaussian noise, idealized *Apparent...* values, determined in turn by

- the true time-averaged pressures: *TrueDiaBP*, *TrueSysBP*, and *TrueMeanBP*;
- the total duration of artifacts within the preceding minute (i.e., the G_{N_j} variables): *BagTime* and *ZeroTime*;
- the average induced pressure to which the transducer is exposed during each event type: *BagPressure* and *ZeroPressure* (these have their own slowly varying dynamics).




Figure 6. ROC curves for the model’s detection of (left) bag events, (middle) zeroing events, and (right) hypertension events.

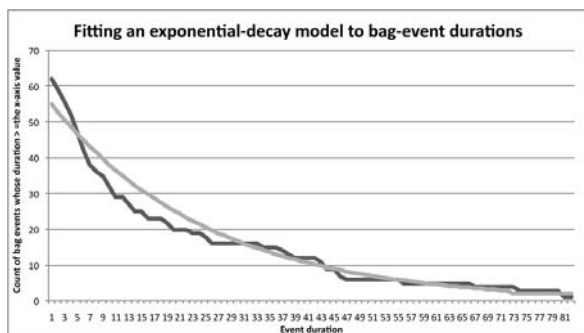


Figure 7. Event count by duration exceeded (dark line), with the fitted exponential (light line).

The *Apparent*... variables are deterministic functions of their parents. For example, we have

$ApparentDiaBP =$

$$BagTime \cdot BagPressure + ZeroTime \cdot ZeroPressure + (1 - BagTime - ZeroTime) \cdot TrueDiaBP$$

The physiological state variables in this model are *TrueSystolicFraction* (the average portion of each heartbeat spent ejecting blood), *TruePulseBP* (the peak-to-trough size of the pressure wave generated by each heartbeat), and *TrueMeanBP*. For simplicity, basic physiologic factors are modeled with random walks weighted toward physiologically sensible values.²

The key event variable in the model, corresponding to f_{Nj} in Figure 4(b), is *EndingValveState*. This has three values for the three possible stopcock positions at the one-minute boundary: open to patient, open to bag, or open to air. The

²More accurate modeling of the physiology actually improves the accuracy of artifact detection, but this point is explored in a separate paper.

CPTs for this variable and for its children (at the next time step) *BagTime* and *ZeroTime* are the ones computed by the method of Section 3. The CPT for *EndingValveState* has $3 \times 3 \times 61 \times 61 = 33,489$ entries.

5. Experimental Results

To estimate the CPT parameters ($P(f_{t+1} = 1 | f_t = 0)$ and $P(f_{t+1} = 1 | f_t = 1)$) for the one-second model, and to evaluate the one-minute model’s performance, we first needed to obtain ground truth for event occurrence. With special dispensation from the ICU nurses, we were able to obtain 300 hours of 1Hz data, a total of more than one million data points. One of us (a physician) examined all of the data and marked the 228 events (119 bag events and 109 zero events). With half the annotated data we estimated the one-second CPT parameters from the event frequencies and durations, verifying that event durations were indeed approximately exponentially distributed (Figure 7). Then, as described in Section 3, we calculated corresponding one-minute-interval CPTs.

After transforming the remaining data into 1-minute averages equivalent to those returned by the regular recording system, we used particle filtering to derive posteriors for true blood pressure and the presence and length of each type of artifact. Figure 6(a) is the ROC curve for the model’s detection of bag events; it reaches a true positive rate of 80% with almost no false positives, or a TPR of 90% with 10% false positives. It does less well with zero-pressure events, as shown in Figure 6(b): a TPR of nearly 70% is achievable with minimal false positives, but beyond that false positives increase quickly.

Detection of hypertensive events is a simple but important clinical application. Standard ICU alarms are threshold-based and susceptible to triggering by artifact, as illustrated

Probabilistic modeling of sensor artifacts in critical care




Figure 8. Two days' blood pressure data for one patient, with the hypertension threshold overlaid. Raw data are on the left; on the right are filtering results showing elimination (here) of false declarations of hypertension.




Figure 9. Sensed blood pressure (dark lines) and inferred true blood pressure (lighter bands, representing mean \pm 1SD) across two observed blood draws. In the upper trace, a blood draw straddles two minutes and is followed a few minutes later by a zeroing; in the lower trace, a blood draw straddles three minutes and is followed by either another short blood draw or a line flush. (The bottom two lines in each tracing show the inferred fraction of each minute occupied by bag or zero artifact.)

in the two days of raw data in Figure 8(left). After filtering with our model, all of the false alarms in this section of data were eliminated, without false negatives, as shown in Figure 8(right). The model's performance on the entire test data set is shown in Figure 6(c). (For the purposes of this ROC curve, the presence or absence at each minute of systolic blood pressure of at least 160mmHg constitutes an event; reference events were defined by a physician during event tagging.) With a true positive rate of 100% in our data set, the model has a false positive rate of less than 10%, a significant achievement given the 86% false-positive rate reported for ICU alarms in general (Tsien & Fackler, 1997).

The model's accuracy in tracking true blood pressure is harder to evaluate since we have no minute-by-minute gold standard. Four of the authors (all physicians) have examined many hours of measured and inferred blood pressure traces, a typical example of which is shown in Figure 9. The results were generally very encouraging and we believe the approach will improve data quality to a level useful for inferring underlying physiological conditions. Where the system's inferences were questionable, it was often the case that examining other sensors helped to reveal whether a pressure change was real or artifactual.

6. Conclusions and Further Work

We have applied dynamic Bayesian network modeling to the problem of handling aggregated data with sub-interval artifacts. In preliminary experiments, this model of a typical blood pressure sensor appears quite successful at tracking true blood pressure and determining the presence, type, and duration of artifacts.

Our approach has reduced the need for learning (as distinct from modeling and inference) to the small but crucial role of determining the distribution of event durations. It would be interesting to test whether direct supervised learning methods could achieve similar results in event detection; it is not obvious, however, how to define the input to a classifier and how to label intervals corrupted by multiple events.

Modified to run at 1Hz, this model could run on-line at the bedside, helping to reduce false alarms. We are currently extending the model to include more sensors and physiological state variables and anticipate further improvements in detection accuracy as a result of combining multiple sensors.

References

Aliferis, C., & Cooper, G. (1996). A structurally and temporally extended Bayesian belief network model: Definitions, properties, and modeling techniques. *Proc. Un-*

certainty in Artificial Intelligence (pp. 28–39).

Beinlich, I., Suermondt, H., Chavez, R., & Cooper, G. (1989). The ALARM monitoring system. *Proc. Second European Conference on Artificial Intelligence in Medicine* (pp. 247–256).

Heldt, T., Long, W., Verghese, G., Szolovits, P., & Mark, R. (2006). Integrating data, models, and reasoning in critical care. *Proceedings of the 28th IEEE EMBS International Conference* (pp. 350–353).

Rao, C. V., & Arkin, A. P. (2003). Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the Gillespie algorithm. *Journal of Chemical Physics*, *118*.

Russell, S. J., & Norvig, P. (2003). *Artificial intelligence: A modern approach*. Upper Saddle River, New Jersey: Prentice-Hall. 2nd edition.

Shachter, R. D. (1986). Evaluating influence diagrams. *Operations Research*, *34*, 871–882.

Tsien, C. L., & Fackler, J. C. (1997). Poor prognosis for existing monitors in the intensive care unit. *Critical Care Medicine*, *25*, 614–619.

Williams, C. K. I., Quinn, J., & McIntosh, N. (2005). Factorial switching Kalman filters for condition monitoring in neonatal intensive care. *NIPS*. Vancouver, Canada.