Using the Blood Pressure Waveform to Reduce Critical False ECG Alarms

GD Clifford,[1] A Aboukhalil,[1] JX Sun,[1], W Zong,[1], BA Janz[1], GB Moody[1], RG Mark[1]

[1]Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

Intensive Care Unit (ICU) false alarm rates can be as high as 86%, leading to a desensitization of the clinical attending staff, slowing of response times and even ignoring true alarms. False alarms are commonly caused by single channel artifacts and could be avoided if information from other independent signals were fused to form a more robust hypothesis of the etiology of the alarm. We used a standard multi-parameter ICU database (PhysioNet's MIMIC DB) to investigate the frequency of false critical (or 'life-threatening') arrhythmia alarms produced by a commercial ICU monitoring system. Multiple expert reviews of the alarms were made using all the relevant files in the MIMIC DB (a total of 21 subjects and 800 hours of waveform data). We found that 25% of the 89 life-threatening alarms were considered false. We then implemented an algorithm to suppress false alarms, using information derived from the arterial blood pressure signal. This simple yet robust strategy was successful in suppressing all false alarms, without suppressing any true alarms.

1. Introduction

False alarms in the Intensive Care Unit (ICU) lead to a disruption of care, impacting both the patient and the clinical staff. ICU alarms produce sound intensities above 80 dB that can lead to sleep deprivation [1, 2] and stress for both patients and staff [3, 4]. More importantly, repetitive false alarms can lead to a desensitization of the clinical staff, confusion, and slowing of response times [1]. In such circumstances, true alarms might be ignored, leading to decreased quality of care [5, 6]. ICU false alarm rates as high as 86% have been reported [7], and between 6% and 40% of ICU alarms have been shown to be true but insignificant. Only 2% to 9% of alarms have been found to be clinically significant [8]. Single channel artifacts are often the cause of false electrocardiogram (ECG) alarms and they could be avoided if information from other independent signals were fused to form a more robust hypothesis of the etiology of the alarm. We hypothesized that the use of a highly correlated signal, such as a pulsatile waveform, to corroborate the alarm category, would allow the suppression of a significant number of false ECG alarms in the ICU. The arterial blood pressure signal (ABP) is perhaps the least noisy pressure signal commonly available, and it is unlikely to contain ECG-related artifacts. In this study, we used a standard multi-parameter ICU database (PhysioNet's MIMIC DB) [9] to investigate the frequency of false *life-threatening* arrhythmia alarms generated by commercial bedside monitors (Hewlett Packard CMS *Merlin*) in a real ICU setting. A system for suppressing the false life-threatening ECG alarms was then developed, using information derived from the ABP waveform.

2. Methods

2.1. Data sources and alarm definitions

A multi-parameter database (MIMIC DB [9]), consisting of over 104 patient-days of real-time signals and accompanying annotations, was chosen for algorithm development. Many of the subjects in this database have simultaneous ECG waveforms, ABP waveforms (sampled at 125Hz with 12-bit resolution) and alarm annotations. The alarms of interest in this project are those classified as life-threatening or *critical* arrhythmia alarms and are defined to be asystole, extreme bradycardia, extreme tachycardia, ventricular tachycardia (VT), and ventricular fibrillation/tachycardia (VF/VT). The definitions of these alarms were established by the manufacturer. The asystole alarm is triggered by a default asystolic pause of 4 sec, that is user-adjustable between 2.5-4 sec. Extreme bradycardia is defined to be a heart rate (HR) below 40 bpm. Extreme tachycardia is defined to be a HR above a default of 140 bpm, adjustable up to 200 bpm for adults. VT is defined as a run of ventricular beats at a rate of at least 100 bpm lasting 5 or more beats. VF/VT is defined as a fibrillatory waveform lasting for at least 3 sec.

We identified a subset of 21 records (totaling 800 monitored hours) that contained simultaneous ECG and ABP waveforms containing a total of 89 critical arrhythmia alarms. Each machine-generated critical alarm was carefully reviewed manually and assessed as true or false by two sequential reviewers. The critical alarm distribution is summarized in Table 1 and discussed in §3.

Table 1. Distribution of the critical arrhythmia alarm types of the MIMIC DB. This dataset is used as the Gold Standard for performance testing of the algorithm.

Alarm Type	True	False	Total
Asystole	7	3	10
Bradycardia	7	4	11
Tachycardia	32	14	46
VT	15	1	16
VF/VT	6	0	6
Total	67	22	89

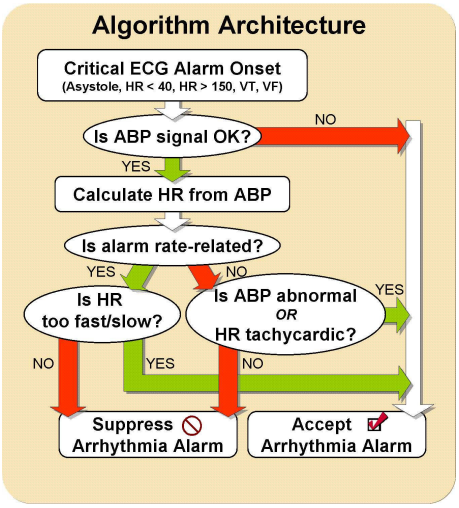

Figure 1. The flowchart outlines the major logical steps of the FA reduction algorithm.

2.2. Algorithm implementation

The proposed algorithm uses ABP waveform information to determine whether an ECG-based critical alarm should be suppressed or accepted. The algorithm may be used as a post-processing module to current ICU monitoring systems to filter critical alarm output in real-time.

The flow of the FA algorithm logic is depicted in Figure 1. At the onset of a critical ECG alarm, the ABP waveform is extracted over a 30 sec analysis window (looking back in time from the onset of the alarm). It should be noted that repeated occurrences of an alarm triggered by a single event were not used in this study, since this would artificially inflate the algorithm's performance. An independent SAI (signal abnormality index) algorithm detailed in [10] is used to determine whether the ABP segment is of high enough quality to yield useful morphological information to adjudicate the alarm. The heart rate is then extracted from the ABP signal and the logic then branches, according to whether the alarm is rate related (bradycardia/tachycardia) or not. If the alarm is not rate related, then further morphological analysis is performed to determine if the alarm is commensurate with the observed ABP waveform. The following sections provide further details concerning each of these processing units.

2.2.1. ABP waveform normality evaluation

A binary beat-by-beat signal abnormality index (SAI) is generated using a set of thresholds and heuristics based on noise level, physiologic ranges and beat to beat variability[10]. We, use the boolean inverse of the SAI, the SNI (signal *normality* index), which has a value of '1' for a normal beat and '0' for an anomalous beat. By computing the mean SNI (mSNI) over a 30 sec analysis window, a measure of the normality of the ABP *segment* is generated, which quantifies the fraction of normal beats within the analysis window. If mSNI $<$ SNIHR (a given threshold < 1), we judge the ABP segment morphology to be 'abnormal' and accept the alarm, since useful information cannot be extracted from the ABP signal. If mSNI $\geq$ SNIHR, further tests can be performed. The subsequent step in the FA algorithm verifies if the alarm is HR-related (asystole, bradycardia or tachycardia) or ventricular-related (VT or VT/VF).

2.2.2. Asystole processing

For an asystole alarm, the ABP waveform is used to compute (a) the longest beat-to-beat interval within the analysis window (in case the asystole resolves itself within the window) and (b) the time between the last beat onset and the end of the analysis window (in case an asystole is sustained beyond the end of the analysis window). If the largest of the latter two values is greater than 3 seconds, an asystole alarm is accepted.

2.2.3. Bradycardia processing

The ABP waveform is used to obtain the 10 largest beat-to-beat intervals within the analysis window. The mean HR is computed based on these 10 candidates. If the mean HR is below the monitor's default threshold, HRMIN $=$ 40 bpm (or any other clinician-adjusted bradycardia HR threshold), a bradycardia alarm is accepted.

2.2.4. Tachycardia processing

The tachycardia processing stage is similar to the one for bradycardia. The 10 shortest beat-to-beat intervals are obtained from the ABP signal. The mean HR is computed based on those 10 candidates. If the calculated mean HR exceeds the monitor's default threshold, HRMAX $=$ 140 bpm (or any other clinician-adjusted tachycardia HR threshold), a tachycardia alarm is accepted.

2.2.5. Ventricular tachycardia and fibrillation/tachycardia

If the alarm is ventricular-related (i.e., VT or VF/VT), the presence of any one of the following two conditions is sufficient to accept the alarm: (a) the ABP segment is not deemed 'normal' at a more stringent level, mSNI $<$ SNIVENT (SNIVENT $>$ SNIHR, since, in comparison to the HR-related alarms, a smaller number of ventricular beats could be corrupted in the analysis window) or (b) the ABP waveform is tachycardic, i.e. the mean HR computed based on the 10 shortest beat-to-beat-intervals is greater than 100 bpm (manufacturer-defined VT threshold). Such design is intended to avoid the possibility of rejecting a true VT alarm where the HR is slightly tachycardic (100-120 bpm) yet displaying an ABP waveform without anomalous beats.

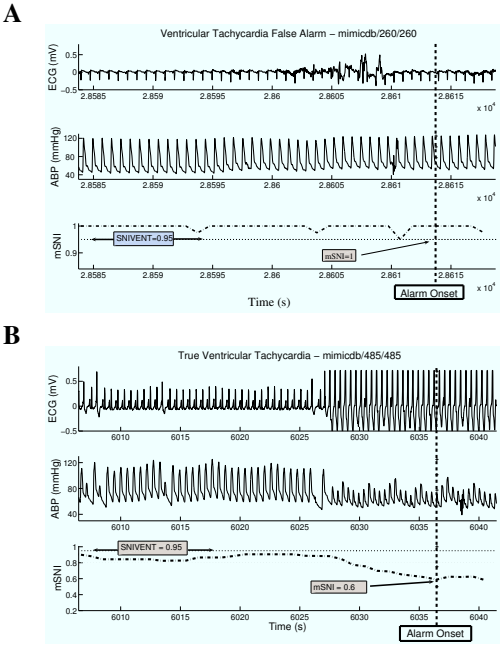

Figure 2. Examples of a false VT alarm suppressed by our algorithm (A), and a true VT alarm that was not suppressed by our algorithm (B). Each plot consists of three time series; an ECG, an ABP waveform and mSNI over 30 sec. The vertical dotted line indicates the alarm onset and the horizontal dotted line indicates the threshold above which the ABP signal is considered to be normal.

Table 2. Summary of FA reduction algorithm performance

	True Alarm	False Alarm
Accepted	67	0
Suppressed	0	22

3. Results

3.1. Human annotation

After expert human annotation of the MIMIC DB life-threatening alarms, we found that 25% of the 89 arrhythmia alarms were considered false (see table 1). Most of the false critical alarms originate from HR disturbances (i.e. asystole, bradycardia and tachycardia), representing approximately one third of the total number of alarms issued by the monitors. The abnormal ventricular-related FA rate of the detected VT and VF/VT events is low, with only one false VT alarm issued (a 6.25% false alarm rate) and with no false VF/VT events.

3.2. Algorithm results

Figure 2 illustrates examples of (A) a false VT alarm suppressed by our algorithm, and (B) a true VT alarm that was not suppressed by our algorithm. Each subplot consists of three time series; an ECG, an ABP waveform and an mSNI (mean SNI over 30 seconds). The dotted vertical line indicates the activation of an alarm by the monitoring equipment, triggered by artifact in Figure 2.A and by true VT in Figure 2.B. Note that in Figure 2.A, the mSNI of the ABP waveform remains high and so we know that it is unlikely that the ECG is manifesting an abnormal (ventricular) rhythm. The alarm is therefore suppressed. In Figure 2.B, mSNI drops below our optimized threshold (mSNI $<$ 0.95), the horizontal dotted line in the third plot. This indicates that there is some abnormality in the ABP waveform either due to noise (so we cannot perform any further tests), or because of the ventricular rhythm. Either way, we cannot suppress the alarm.

The algorithm's detection and rejection performance was compared to the 'Gold Standard' annotations, and the results were tabulated in Table 2. The algorithm successfully rejects all false alarms (22/22) and correctly accepts all true alarms (67/67), yielding a sensitivity of 100% and a positive predictivity of 100%.

3.3. Sensitivity analysis

Given a relatively humble dataset of 89 alarms, it was not deemed viable to separate the data into development and test subsets. Instead, a sensitivity analysis was conducted to quantify the robustness of our algorithm's performance to changes in its threshold values. This also helped to ensure that the algorithm's performance was not substantially dependent on the particular threshold values, and to demonstrate that the thresholds were not over-tuned to this data set.

The algorithm's operating point is SNIHR = 0.8, SNIVENT = 0.95, and analysis window length = 30 sec. In

the analysis, each of the 3 thresholds (SNIHR, SNIVENT, and analysis window length) was individually varied one at a time around the operating point while maintaining the other 2 constant.

The results of individually varying the SNIHR and SNIVENT values from 0 to 1, and the analysis window from 10 to 120 seconds are now summarized. Correct detection with a sensitivity and positive predictivity of 100% (i.e. FP = 0, FN = 0, and TP rate = 100%) occurs for an SNIHR ranging between 0.70-0.86, an SNIVENT between 0.92-1.00, and analysis window length between 25-45 sec. At low SNIHR values (< 0.70), the signal's morphology is more readily considered abnormal. Therefore, the algorithm becomes more stringent and suppresses more true alarms, thereby becoming extremely unsafe. On the other hand, when SNIHR exceeds 0.86, the algorithm accepts alarms more easily and increases the chances of issuing a false alarm. At SNIVENT values below 0.92, the ABP waveform's morphology is more readily regarded as abnormal and more true alarms are hence suppressed. Finally, by varying the analysis window length, performance degrades both below 25 and above 45 sec; more true alarms are suppressed and more false alarms are accepted. A shorter time window presents less data to make an accurate computation, whereas a longer window includes more data that could be unrelated to the alarm. Therefore, a time window of 30 sec was found to be a reasonable trade off in this case.

4. Conclusions

Our results on human annotations of false alarms in the MIMIC DB indicate that 25% of the critical ECG arrhythmia alarms (asystole, bradycardia, tachycardia, VT, and VF/VT) are false. Our subsequently constructed FA suppression algorithm uses information from the ABP waveform to suppress critical ECG arrhythmia FAs. Our algorithm compares favorably with the only other system that has been publicly reported: GE Medical's *Intellirate*, which has been reported to reduce FAs by up to 50%, with false asystole alarm suppression as high as 89%. Our FA suppression algorithm presented in this paper has been shown to suppress 100% of life-threatening false ECG arrhythmia alarms in the ICU, over a wide range of thresholds, on a publicly available database of 800 hours of ICU data. Moreover, the algorithm did not suppress any true alarms.

The human false/true alarm annotations, together with the SNI values for each beat derived in this study will be made publicly available through PhysioNet [11]. A limitation of this pilot study is the fact that algorithm development and evaluation were conducted using the same data. Follow-on work will utilize the much larger MIMICII DB [12], which has an order of magnitude more data and alarms. Analysis of this new data will allow the testing and development of this algorithm on a more varied and representative set of waveform examples.

Acknowledgments

The authors gratefully acknowledge support from the NIH/NBIB grant R01 EB001659.

References

[1] Chambrin M. Review: Alarms in the intensive care unit: how can the number of false alarms be reduced? Critical Care May 2001;5(4):184–188.

[2] Meyer T, Eveloff S, Bauer M, Schwartz W, Hill N, Millman R. Adverse environmental conditions in the respiratory and medical ICU settings. Chest 1994;105(4):1211–1216.

[3] Aaron J, Carlisle C, Carskadon M, Meyer T, Hill N, Millman R. Environmental noise as a cause of sleep disruption in an intermediate respiratory care unit. Sleep 1996; 19:707–710.

[4] Cropp A, Woods L, Raney D, Bredle D. Name that tone. The proliferation of alarms in the intensive care unit. Chest 1994;105(4):1217–1220.

[5] Donchin Y, Jacob S. The hostile environment of the intensive care unit. Current Opinion in Critical Care August 2002;8(4):316–320.

[6] Imhoff M, Kuhls S. Alarm Algorithms in Critical Care Monitoring. Anesth Analg 2006;102(5):1525–1537.

[7] Lawless S. Crying wolf: false alarms in a pediatric intensive care unit. Crit Care Med Jun 1996;22(6):981–985.

[8] Tsien C, Fackler J. Poor Prognosis for Existing Monitors in the Intensive Care Unit. Crit Care Med Apr 1997; 25(4):614–619.

[9] Moody GB, Mark RG. A database to support development and evaluation of intelligent intensive care monitoring. Computers in Cardiology 1996;23:657–660.

[10] Sun J, Reisner A, Mark R. A signal abnormality index for arterial blood pressure waveforms. Computers in Cardiology 1999;33.

[11] Mark R, Moody G. PhysioNet: the research resource for complex physiologic signals: MIMIC; A Multiparameter Database. http://www.physionet.org/physiobank/database/mimicdb/.

[12] Saeed M, Lieu C, Raber G, Mark R. MIMIC II: A massive temporal ICU patient database to support research in intelligent patient monitoring. In Comput. Cardiol. 2002. Los Alamitos: IEEE Computer Society Press, 2002; 641–644.

Address for correspondence:

Gari D. Clifford
Harvard-MIT Division of Health Sciences & Technology
45 Carleton St.,
Cambridge MA 02139, / USA
gari@mit.edu
http://alum.mit.edu/www/gari/