Journal of Biomedical Informatics 45 (2012) 913–921

Contents lists available at SciVerse ScienceDirect

Journal of Biomedical Informatics

journal homepage: www.elsevier.com/locate/yjbin

Predictive combinations of monitor alarms preceding in-hospital code blue events

Xiao Hu [a,b,c,*], Monica Sapo [c], Val Nenov [c], Tod Barry [e], Sunghan Kim [a], Duc H. Do [d], Noel Boyle [d], Neil Martin [c]

[a] Neural Systems and Dynamics Laboratory, Department of Neurosurgery, David Geffen School of Medicine, University of California, Los Angeles, United States
[b] Biomedical Engineering Graduate Program, Henry Samueli School of Engineering and Applied Science, University of California, Los Angeles, United States
[c] Department of Neurosurgery, David Geffen School of Medicine, University of California, Los Angeles, United States
[d] UCLA Cardiac Arrhythmia Center, David Geffen School of Medicine, University of California, Los Angeles, United States
[e] Quality Management Service, UCLA Ronald Reagan University Medical Center, Los Angeles, United States

ARTICLE INFO

Article history:
Received 4 October 2011
Accepted 9 March 2012
Available online 24 March 2012

Keywords:
Association rule mining
Alarms
Cardiac arrest
Code blue
Alarm fatigue

ABSTRACT

Bedside monitors are ubiquitous in acute care units of modern healthcare enterprises. However, they have been criticized for generating an excessive number of false positive alarms causing alarm fatigue among care givers and potentially compromising patient safety. We hypothesize that combinations of regular monitor alarms denoted as SuperAlarm set may be more indicative of ongoing patient deteriorations and hence predictive of in-hospital code blue events. The present work develops and assesses an alarm mining approach based on finding frequent combinations of single alarms that are also specific to code blue events to compose a SuperAlarm set. We use 4-way analysis of variance (ANOVA) to investigate the influence of four algorithm parameters on the performance of the data mining approach. The results are obtained from millions of monitor alarms from a cohort of 223 adult code blue and 1768 control patients using a multiple 10-fold cross-validation experiment setup. Using the optimal setting of parameters determined in the cross-validation experiment, final SuperAlarm sets are mined from the training data and used on an independent test data set to simulate running a SuperAlarm set against live regular monitor alarms. The ANOVA shows that the content of a SuperAlarm set is influenced by a subset of key algorithm parameters. Simulation of the extracted SuperAlarm set shows that it can predict code blue events one hour ahead with sensitivity between 66.7% and 90.9% while producing false SuperAlarms for control patients that account for between 2.2% and 11.2% of regular monitor alarms depending on user-supplied acceptable false positive rate. We conclude that even though the present work is still preliminary due to the usage of a moderately-sized database to test our hypothesis it represents an effort to develop algorithms to alleviate the alarm fatigue issue in a unique way.

1. Introduction

Bedside monitors are ubiquitous in acute care units of modern hospitals. However, they are often criticized for generating an excessive number of false positive and false alarms [1–7]. Frequent false positive alarms not only create annoying distractions but also can cause alarm fatigue for bedside care givers so that attentions to critical alarms are missed raising serious patient safety concerns [8–12]. Indeed, recent mainstream reports have published cases of avoidable patient deaths that were unfortunately related to the alarm fatigue/desentization among bedside care givers [8]. Therefore, it is imperative to investigate different strategies to improve patient monitor alarm generation and management.

The issue of false alarms and false positive alarms has been well studied. In a recent report [6], only 15% of alarms have been found to be clinically relevant in a medical intensive care unit (ICU). In an emergency room setting, it has been reported that only 0.7% of alarms are true positives meaning that they have detected adverse events and led to clinical interventions [5]. Similar findings regarding a high percentage of clinically irrelevant alarms have been reported in diverse ICU environments [1–4]. False positive alarms can be caused either by false alarms due to noise and artifacts in signals or by inappropriate alarming criteria that are too generic and sensitive. Indeed, most of the threshold-based alarms despite being true alarms are false positives [13]. Extensive research efforts have been put into developing solutions to reduce the false positive rate of monitor alarms [13]. Understandably, majority of these efforts have been targeted at improving signal processing aspects of alarm generation with the hope that robust signal processing can lead to fewer false alarms [14–19]. Reducing the false positive rate beyond reducing the number of false alarms is more

* Corresponding author. Address: 18-265 Semel, Box 703919, 10833 Le Conte Avenue, Los Angeles, CA 90095, United States. Fax: +1 310 794 2147.
E-mail address: xhu@mednet.ucla.edu (X. Hu).

1532-0464/$ - see front matter © 2012 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jbi.2012.03.001

challenging because of the need for highly sensitive monitoring in an acute care setting. Advanced pattern recognition of biomedical signals has also been advocated as a method to create intelligent alarms that are hopefully more specific without sacrificing sensitivity [20–24]. However, these studies are in research phase and there is a demand to create annotated databases to evaluate different intelligent alarm algorithms [25].

In the present work, we propose a different strategy to achieve a better alarm generation process by directly analyzing raw alarm messages from the monitors. A direct analysis of alarms has been undertaken in existing studies but the focus has been on annotating individual alarms by trained observers to categorize them into false and true positive alarms [1–6]. This effort indeed matches the prevailing patient monitoring practice where care givers process alarms one by one as they go off. Little time is available for them to recall historical alarms and then manually associate them with the current alarm to create a more holistic assessment of patients. What is missed in this single-alarm practice is the ability to account for potential predictive patterns arising from a combination of different single alarms. Therefore, we are proposing and testing a method that is capable of mining a collection of monitor alarms to search for frequent but also specific combinations of encoded monitor alarms to predict certain adverse event. We have chosen this event to be in-hospital code blue arrests but the developed algorithm works not only for this choice of target events.

To the best of our knowledge, the present work is probably the first study in terms of directly mining patient monitor alarms. However, some methods used here have been well studied in other fields. A key technique is to find frequent itemsets given a collection of alarm sequences from many patients. This technique is part of a well known association rule mining algorithm [26]. Similar techniques have been used in fault diagnosis to automate the process of processing logs of computer-generated systems errors to predict impending fault of a complex computer system [27]. Discretization of continuous variables is used in our method as well, which is also a well studied topic [28,29]. Apart from these existing methods that our proposed algorithm has leveraged, there are additional and novel elements in our proposed algorithm that will be the focus in next section.

2. Materials and methods

2.1. A frequent itemset based alarm mining algorithm

Table 1 shows the composition of monitor alarms by using four examples. A raw monitor alarm includes a unique alarm code assigned by the monitor manufactory, a textual label of the alarm which is often uniquely mapped to the alarm code, an optional polarity indicator that denotes whether a parametric alarm exceeds an upper bound (HI) or a lower bound (LO) threshold, an optional value at which the preset alarm thresholds have been crossed to trigger this alarm, and the timestamp when this alarm occurs. There are four built-in levels of alarms as determined by bedside monitors, which are usually set up by a unit-based policy. These four levels are: crisis alarm, patient advisory alarm, patient warning alarm, and system warning alarm.

Fig. 1 displays the flowchart of the proposed alarm data mining algorithm. To facilitate the subsequent discussion, a combination of encoded raw alarms that co-occur within a temporal window is termed a SuperAlarm pattern. The goal of our algorithm is to construct a set of predictive SuperAlarm patterns from two collections of raw alarm data. As depicted in this flowchart, the first collection (cases) consists of alarms that precede code blue events in multiple patients. The second collection (controls) consists of alarms from a set of control patients. These two collections go through two different branches of processing. The case data are used to find SuperAlarm patterns occurring frequently within a T_w-long window that immediately precedes code blue events. The control data are used to filter out those SuperAlarm patterns identified for code blue patients that have also occurred frequently for control patients. This is achieved by sampling alarms for control patients in consecutive windows of n-hour long starting from the beginning of the monitoring to the end. Alarms thus samples are assumed to be representative of the whole course of patient monitoring. Within each window, alarms are sampled from a randomly placed T_w-long segment. In this way, a false positive rate can be readily computed for each SuperAlarm pattern during the training phase. The algorithm also needs to discretize the value field for parametric alarms. The discretization algorithm [30] requires the supply of both case and control data. Hence, alarms within the first 2 h of monitoring are used to generate the discretization schema.

Having presented the general idea behind the mining algorithm, we describe the individual processing blocks of this algorithm in the following subsections.

2.1.1. Alarm pre-processing

Due to the fact that a bedside monitor can have multiple input ports to accommodate multiple monitoring modalities, the same device can be plugged into any of those ports and results in different labeling of the same alarm. In the example shown in Table 1, the arterial line (A-line) was plugged into port #1 and hence ART1 is part of the label. The first pre-processing task is therefore to make an alarm agnostic to the specified port number. In the second pre-processing task, we treat the alarms from noninvasive devices as equivalent to those from its invasive counterpart and hence alarms from invasive and noninvasive blood pressure are merged. Considering that the value of a measurement that triggers an alarm can be good indicator of the severity of the alarm, we use a discretization algorithm to further divide a given alarm with value field into sub-codes. We call such a process a regular alarm encoding. We have used a data-driven approach class-attribute contingency coefficient (CACC) [30]. This approach needs both case data and control data to create a two-class discrimination problem to find the optimal discretization that will result in the best corre-

Table 1
Illustration of the composition of monitor alarms using four example alarms. Polarity is an indicator that denotes whether a parametric alarm exceeds an upper bound (HI) or a lower bound (LO) threshold.

Alarm Code	Label	Level	Polarity	Value	Timestamp
90	ART1 S	Patient warning	LO	80	6/23/2011 14:50:11
Systolic arterial blood pressure at port #1 crosses the preset alarm lower bound at a value of 80 mm Hg at the specified time point.					
89	ART1 S	Patient warning	HI	180	6/2/2011 4:30:11
Systolic arterial blood pressure at port #1 crosses the preset alarm upper bound at a value of 180 mm Hg at the specified time point.					
1	Asystole	Crisis	N/A	N/A	6/23/2011 5:20:10
An asystole alarm is trigged at the specified time point					
190	NBP S	Patient warning	HI	160	6/23/2011 11:50:11
Noninvasive systolic arterial blood pressure crosses the preset alarm upper bound at a value of 160 mm Hg at the specified time point					