

Hypotension Risk Prediction via Sequential Contrast Patterns of ICU Blood Pressure

Shameek Ghosh, Mengling Feng, Hung Nguyen, *Senior Member, IEEE* and Jinyan Li

Abstract—Acute hypotension is a significant risk factor for in-hospital mortality at intensive care units (ICUs). Prolonged hypotension can cause tissue hypoperfusion, leading to cellular dysfunction and severe injuries to multiple organs. Prompt medical interventions are thus extremely important for dealing with acute hypotensive episodes (AHE). Population level prognostic scoring systems for risk stratification of patients are suboptimal in such scenarios. However, the design of an efficient risk prediction system can significantly help in the identification of critical care patients, who are at risk of developing an AHE within a future time span. Towards this objective, a pattern mining algorithm is employed to extract informative sequential contrast patterns from hemodynamic data, for the prediction of hypotensive episodes. The hypotensive and normotensive patient groups are extracted from the MIMIC-II critical care research database, following an appropriate clinical inclusion criteria. The proposed method consists of a data preprocessing step to convert the blood pressure time series into symbolic sequences, using a symbolic aggregate approximation algorithm. Then, distinguishing subsequences are identified using the sequential contrast mining algorithm. These subsequences are used to predict the occurrence of an AHE in a future time window separated by a user-defined gap interval. Results indicate that the method performs well in terms of the prediction performance as well as in the generation of sequential patterns of clinical significance. Hence, the novelty of sequential patterns is in their usefulness as potential physiological biomarkers for building optimal patient risk stratification systems and for further clinical investigation of interesting patterns in critical care patients.

Index Terms—Acute hypotension, Mean arterial pressure, Symbolic sequences, Sequential pattern mining

I. INTRODUCTION

IN the past few years, there has been a significant rise in patient monitoring devices aggregating large-scale patient data in intensive care units. Typically, most of this huge volume of data has remained underutilized, leading to slower progress in medical research. However, with increasing demand on healthcare organizations, there is now an urgent necessity to provide improved access and quality of care

at lesser costs. As evidence obtained from modern data-driven techniques have contributed to significant advances in critical care patient diagnosis, such efforts have resulted in an improved understanding of diseases and guided appropriate medical interventions.

Appropriate clinical diagnosis of impending critical events is extremely important in an ICU, since rapid physiological changes cause critical patient instabilities that require immediate medical interventions. Conventional early warning monitoring systems turn out to be suboptimal in such cases. Existing systems embed a set of predefined clinical rules, which act on vital signs data, to raise an alarm reactively. Moreover, they are also known to generate a significant number of false alarms in ICUs [2]. In addition, the current systems do not account for the dynamic nature of complex physiological processes in a given time period. Hence, there exists a need for predictive technologies, which can act proactively for advanced medical decision-making in critical care units.

Hemodynamic monitoring is an essential mechanism in ICUs generating a significant amount of streaming blood pressure (BP) data. Acute hypotensive episodes (AHE) are defined as a sudden drop of patient blood pressure spanning over an extended time period. An AHE can lead to decreased tissue perfusion, which in turn can be a cause of multiple organ damages. Hemodynamic instabilities can be life-threatening to the concerned patients. On the other hand, if such instabilities are detected ahead of time to limit the effects of a life threatening event, then there are significant benefits associated with the outcomes.

The effectiveness of medical outcomes is generally assessed by the risk of mortality and also involves the costs of treatment. For critical care patients, these factors tend to rise with time. Thus, the effectiveness of individual medical outcomes is strongly dependent on well-informed patient interventions. Proactive interventions are staged on the basis of clinical evidence of impending events. Such evidence needs to have two significant characteristics viz. predictive capability and clinical interpretability. The importance of clinical interpretability stems from the requirement of a clinician's enhanced degree of understanding of the patient's physiological condition. Such knowledge is fundamental for the selection of an optimal treatment plan.

A knowledge discovery based predictive system can meet this demand. Usually, such a predictive system takes into account time-based micro physiological events during a patient's ICU stay. It is able to make significant associations of interpretable clinical evidence to future hemodynamic behaviour. Accordingly, it has a strong potential for a reduction

S. Ghosh is with the Advanced Analytics Institute, Faculty of Engineering and IT, University of Technology Sydney (UTS), Australia. E-mail: (Shameek.Ghosh@student.uts.edu.au).

M. Feng is with the Laboratory for Computational Physiology, Harvard-MIT Joint Health, Science and Technology, Massachusetts Institute of Technology, Cambridge, USA and Institute of Infocomm Research, Singapore. E-mail: (mfeng@MIT.EDU).

H. Nguyen is with the Centre for Health Technologies, Faculty of Engineering and IT, University of Technology Sydney (UTS), Australia. E-mail: (hung.nguyen@uts.edu.au).

J. Li is with the Advanced Analytics Institute, Faculty of Engineering and IT, University of Technology Sydney (UTS), Australia. E-mail: (Jinyan.Li@uts.edu.au). Corresponding author.

in operational costs, increase in efficiency, the development of novel goal directed treatments and scheduling of additional ICU services.

The aim of this study is to identify discriminative hemodynamic sequential patterns via a novel data mining method for the risk stratification of ICU patients. These patterns are later utilized to distinguish hypotensive episodes from normotensive cases.

The informative sequential patterns are extracted from a large-scale patient population in the MIMIC-II critical care research database [3]. The MIMIC-II (Multiparameter Intelligent Monitoring in Intensive Care) database is a publicly available critical care data resource, encompassing a diverse and large population of ICU patients over the last 10 years. It comprises of high resolution temporal data including lab results, discharge notes, physiological trends and waveforms. The database has been widely used to support numerous research studies in the fields of epidemiology, clinical decision-rule improvement, and ICU alarm systems.

One important novelty of the current study is the application of a sequential contrast pattern mining strategy in the extraction of clinical episodes of arbitrary length, which are a characteristic of specific critical conditions like an AHE. The present study can thus meet the need to generate novel medical insights from the data of intensive care units and discover clinically relevant episodes separated by time windows.

Overall, our contributions made by this study include: (i) the application of a contrast pattern mining technique in the field of critical care informatics, (ii) a new method for generating predictive alerts for hypotensive episodes in an ICU, and (iii) validation of our method on data extracted from a large-scale de-identified critical care research database like the MIMIC-II.

II. PROBLEM DEFINITION

Acute hypotension is a clinical symptom showing a significant drop in mean arterial pressure (MAP) values for extended periods of time. The mean arterial pressure is often used in medicine as a popular measure of blood pressure, which can be derived from the systolic (SP) and diastolic pressure (DP) as given by equation (1).

$$MAP = [2(DP) + SP]/3 \quad (1)$$




Fig. 1. Acute Hypotensive Episode over a time period exceeding 30 minutes, when $MAP \leq 60$ mmHg

Although hypotension is not categorized as a disease state, it is considered to be a frequent ailment among the general population and especially among females. Owens et al [28] reported

a prevalence of 49% hypotensive patients in a prevalence study of a general population cohort. Existing studies have indicated that hypotension is associated with morbidities stemming from dizziness and fatigue. Hypotensive subjects have previously demonstrated lower blood pressure, along with lower weights and had lesser likelihood of a family history of vascular disease or hypertension. However, in a diagnostic setting, actual prevalence can be dependent upon associated stress, anti-hypertensive medications and diuretics [23]. Neurological diseases are also associated with an increasing likelihood of AHEs in an ICU. Depending on various definitions of hypotension, MAP values falling below the threshold range of 60-80mmHg for 30 minutes, could trigger an acute hypotensive episode. Figure 1 illustrates such a scenario, where MAP values sustain below 60 mmHg for a time period ≥ 30 minutes.

A. A Formulation of the AHE Prediction Problem

Numerous studies report that hypotension could lead to critical events like acute kidney injury, severe sepsis, acute coronary syndrome and shock [46], [47], [48], [21]. To enable prompt interventions, it is therefore important to predict an AHE ahead of time. Predicting an AHE can be formulated as a problem of classification of an admitted patient's mean arterial pressure into a hypotensive or normotensive regime. The prediction of the mean arterial pressure in a future time window is central to the current study. An illustration of the AHE prediction problem is provided in Figure 2.

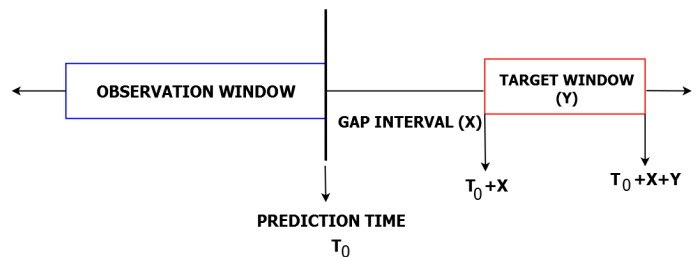


Fig. 2. Observation and Target Windows with a Time Gap Interval

According to Figure 2, a user-defined MAP time series observation window of length 30 or 60 minutes, is provided as historical data. The time series observation window is subsequently utilized to predict the given MAP's class (hypotensive or normotensive) in a future target window of 30 minutes. Moreover, the observation and the target windows are separated by an user-defined gap interval of 60 and 120 minutes. The problem can be interpreted as that of performing an AHE prediction in a future time window, given the MAP observation data and a gap interval of one or two hours between the observation and the forecasting time windows.

B. Related Work

There have been a number of studies using pattern recognition techniques for the analysis of hypotensive behaviour. Wavelet-based similarity measures from blood pressure time series had been proposed to predict vasopressor onsets [6]. Ghaffari et al [7] have demonstrated the use of Hilbert-transform based techniques for predicting AHEs. In 2009,

the Physionet AHE prediction challenge was instituted to advance the development of state-of-the-art techniques [5], including neural networks, support vector machines and numerous statistical indices as features, for the prediction of AHE [8], [9], [12]. In some of these cases, historical time windows used for observations were considered as five minutes prior to the onset of an AHE. Accordingly, Wang et al [13] have reported that medical pattern extraction was particularly challenging, owing to their longitudinal and sparse nature. While addressing this concern, Syed et al [14] reported the development of motif mining methods, which were tested on long-range cardiovascular time series datasets. Moreover, Lee and Mark [15] reported the extraction of hemodynamic patterns for hypotension through artificial neural networks. Recently, Ghassemi et al [51] demonstrated the application of global optimization techniques for principal extraction of dynamical features, which are more predictive of patient outcomes than severity of illness scores accounting for multiple non-dynamic features of the patient.

For the area of predictive data mining for monitoring applications, previous research has reported the development of numerous pattern mining techniques. Typically, existing research tends to identify problems in either of two directions viz. short-term predictive modelling with the objective of generating daily alerts for physicians or long-term predictive modelling aimed at population level prognosis [22]. Monitoring systems help in capturing signals that can be used to identify time varying phenomenon, instead of traditional generation of alerts, which are known to generate a lot of false reports [24]. To overcome this weakness, intelligent noise removal methods are used as low pass filters which can aggregate high resolution signal frames and a number of good measurements [25]. The processed input is then used for classification and regression problems, although the concerned method may or may not consider temporal aspects of the data. A wide range of ICU prediction tasks focus on the extraction of statistical features from medical time series and making them time-independent. For example, impending ICU physiological instability has been predicted by decision rules from time series data [26]. Bayesian artificial neural networks were employed for observation windows 15 and 30 minutes before hypotension for traumatic brain injury subjects resulting in 86% specificity and 41% sensitivity [27]. Second day ICU discharges were predicted by gaussian processes [30]. Celi et al [29] also employed Bayesian networks to predict day-two fluid requirements for the study of patient inflammatory responses.

Apart from final prediction outcomes, medical decision makers also expect to discover insights relating to the processes employed on longitudinal patient records. Research on such data, begins with complex data transformation procedures by developing temporal abstractions to represent temporal relations between time intervals. Previous studies have reported the extraction of meaningful temporal patterns from a diabetes dataset [31], [33], [38]. **Prior to this, Tseng and Lee [54], [55] had reported temporal pattern-based classifiers for effective classification by sequences for atrial fibrillation datasets.** Additionally temporal patterns were also used to

predict the hospitalization of hemodialysis patients [35], [36]. A time-series knowledge mining method was used to discover frequent temporal patterns for patients who required mechanical ventilation for greater than 24 hours [37]. Toma et al [34] utilized frequent temporal patterns to capture the evolution of organ failures status in a set of patients. Temporal history of patient event codes have also been reportedly used for mining frequent sequences of events to understand various illnesses [32]. **In this context, Perer et al [52] demonstrated the clinical usefulness of frequent temporal sequences by an interactive and visual analytics platform for mining sequences of ICD-9 codes to understand disease progression. Similar visual analytics platforms have been shown to have a greater clinical importance in the mining of medical event sequences having strong associations with specific disease outcomes [53].**

III. METHODOLOGY

The identification of sequential patterns is related to detecting subsequences contained within training sequences. According to the problem constraints, well-defined representative patterns may be grown, which display strong support in the concerned training sequences. Sequential pattern mining strategies can provide a useful alternative to mining interesting patterns of physiological time series data, which may help discover significant insights in the form of important clinical episodes. In the following sections, we describe the various stages viz. data extraction, pre-processing and mining of sequential contrast patterns that are over-represented in the hypotensive training samples and under-represented in the normotensive samples.

A. Data Extraction

The data of the study is a relevant subset of the MIMIC-II database, using a suitable data inclusion/exclusion criteria [15]. The MIMIC II is a large-scale intensive care unit database consisting of more than 30,000 patients with numerous patient variables, aggregated from patient health care records and physiological waveforms over a period of more than 10 years. The physiological time series waveforms data are organized into records, identified using unique patient identifiers. A specific patient identifier may correspond to multiple ICU stays. Thus, time series data for each ICU stay maintains a unique ICU stay identifier. The extracted subset of records are also satisfied the following conditions, before extraction.

- The record had to be of an adult patient.
- Each patient time series constituted of minute-by-minute numeric samples, for at least the mean arterial blood pressure.
- Corresponding clinical records existed for the waveform records in MIMIC II.

As recommended by Lee and Mark [15], we considered the following inclusion criteria, while compiling the data examples. As described in Figure 1, each data sample comprised of three time intervals as follows.

- a 30 or 60 minutes MAP observation window

- a 30 minutes target window
- a time interval gap of 60 or 120 minutes, which separates the observation and target windows.
- There exist seven categories for the ICD-9 code for hypotension (458.0 - 458.9) as shown in **Table I**. Hypotensive records were selected by pattern matching over the higher level numerical classification of 458 in MIMIC-II.

TABLE I
ICD-9 CLASSIFICATION OF HYPOTENSION

ICD-9 Code	Disease
458.0	Orthostatic hypotension
458.1	Chronic hypotension
458.2	Iatrogenic hypotension
458.21	Hypotension of hemodialysis
458.29	Other iatrogenic hypotension
458.8	Other specified hypotension
458.9	Hypotension unspecified

A target window was labelled either as normotensive (control) or hypotensive. The labelling of a target window as hypotensive (HE) was subject to satisfying a 30 minute period of time for which MAP was less than 60 mmHg and greater than 10 mmhg, for 90% of the time period. In contrast, a 30 minute window which did not satisfy the given HE definition as above was labelled as a normotensive (control) sample. Moreover, corresponding to each target window, the extracted MAP observation windows were also verified to be within the 10-200 mmHg range.

Two data extraction mechanisms were considered viz. single and multiple modes. For single mode compilation, a single hypotensive or normotensive example was constructed from each separate patient waveform record. On the other hand, the multiple compilation mode considered a sliding window of 30 minutes, and all those examples were constructed, whenever satisfying the conditions for the observation and target windows.

In addition to the datasets extracted using the given inclusion criteria, hypotensive and normotensive datasets were also employed from the Physionet 2009 challenge [5]. For the challenge datasets, their MIMIC II waveform signals were divided into two groups viz. H (hypotensive) and C (control) respectively. The groups H and C were further subdivided into H1, H2 and C1, C2. Each sub-group were defined to have the following properties.

- H1: Patients receiving pressor medication.
- H2: Patients not receiving pressor medication.
- C1: Patients with no acute hypotensive episodes during entire hospital stay.
- C2: Patients having AHE before or after the forecast window.

Accordingly, two challenge prediction tasks were constituted as follows.

- Event I: Patient risk classification between H1 and C1
- Event II: Patient risk classification between H and C

Moody and Lehman [5] reported that the groups H1 and C1 indicated the extremes of AHE-associated risks.

B. Data Discretization

Physiological data often comprise of repetitive elements. To identify interesting patterns, a natural extension is to transform the real-valued physiological time series into string representations for mining symbolic discrete patterns [14]. Subsequently, we employed the symbolic aggregate approximation method [16] to segment the original MAP signal into discrete intervals and assigned an alphabetic label to each discrete region. This process transforms the continuous MAP data into a symbolic sequence, and enables the use of numerous pattern mining algorithms. The symbolic aggregate approximation (SAX) technique has emerged as a popular and efficient technique, producing an informative symbolization of large-scale time series data. Typically, SAX converts the continuous time series into a piecewise aggregate approximation (PAA) form [16]. Later, the PAA series is converted to a symbolic sequence. Each MAP time series, before being discretized, undergoes a normalization process having a mean of 0 and variance 1. The SAX strategy selects breakpoints using a gaussian distribution, such that the discrete symbols are equiprobable in the time series. For example, to transform a normalized time series using five symbols, the discrete regions are specified by $[-\infty, -0.84, -0.25, +0.25, +0.84, +\infty]$. The symbolic representation adopted by SAX characterizes the inherent properties of the time series data. Consequently, an equiprobable distribution of symbols is maintained in the given time series [16].

In the process, SAX provides an effective discretization platform, which can be utilized to create efficient pattern mining and indexing algorithms for medical purposes. Figure 3 illustrates a visual representation of a real-valued time series being converted to a symbolic form, using four symbolic regions.




Fig. 3. Discretization by Symbolic Aggregate Approximation using 4 symbols

C. Mining gap-constrained sequential contrast patterns

In studies related to binary or multi-class classification, the central objective is to develop a prediction model, which is capable of distinguishing an incoming signal using its inherent properties and assign a target label as the predicted outcome. Typically, in data mining problems, there exists a strong motivation to discover differentiable patterns' characteristic of disparate groups of data, that are used for prediction of records. Mining emerging patterns from distinctively labelled groups of relational data was initially introduced by Dong and Li [17]. However, the immediate application of emerging patterns to sequence databases was not possible owing to ordering of a sequence, and due to multiple occurrences of items in a sequence. Accordingly, the concept of emerging

substrings was suggested [42]. Substrings are a special case of subsequences, where each consecutive symbol is separated by a gap interval of 0. Yet, an important aspect to note is that significant sequential episodes may not manifest as consecutive symbols existing in time-series symbolic sequences of interest. Thus, the identification of episodes having significant events ordered sequentially, while having arbitrary gap intervals between events, can be extremely useful. Towards this purpose, a number of algorithms have been reported [43]. In the present study, we intend to discover gap-constrained contrast subsequences from disparate groups of sequence data, using the principles of frequency support. In the following sections, the various definitions and processes associated with the extraction of gap-constrained sequential contrast patterns are described.

1) *Sequential Patterns*: Let there be a set of distinct items denoted as I . I can also be called the alphabet set and $|I|$ is the size of the alphabet set. A sequence S defined over I may be denoted as $e_1 - e_2 - \dots - e_n$, such that e_i belongs to I for $1 \leq i \leq n$. Accordingly, we consider univariate sequences where e_i represents a single item from I . A sequence $S' = e_{i_1} - e_{i_2} - \dots - e_{i_m}$ is said to be contained in a sequence $S = e_1 - e_2 - e_3 - \dots - e_n$, such that $1 \leq i_1 \leq i_2 \leq \dots \leq i_m \leq n$. For example, a subsequence CD is contained in $CAAD$, but not DC . Hence, the order of the sequence S' is maintained in S , although items in S' are not consecutive in S . This indicates the existence of gap intervals between the items of S' .

Definition 3.3.1.1: (Max-Prefix) The max-prefix of the sequence $S = e_1 - e_2 - \dots - e_k$ is given by $e_1 - e_2 - \dots - e_{k-1}$. It constitutes the leading sequence of elements in S , without the final item of S .

Definition 3.3.1.2: (Occurrence of a Subsequence) Given the sequences, $S = e_1 - e_2 - \dots - e_n$ and $S' = e_{i_1} - e_{i_2} - \dots - e_{i_m}$, S' occurs in S if $1 \leq i_k \leq n$ and $e_k = e_{i_k}$ for all $1 \leq k \leq m$, and $i_k \leq i_{k+1}$ for $1 \leq k \leq m$. For example, given sequences $S = XZXYZZY$ and subsequence $S' = XY$, there are four occurrences of S' in S at the positions - $\{1, 5\}$, $\{1, 7\}$, $\{3, 5\}$ and $\{3, 7\}$.

Definition 3.3.1.3: (Satisfaction of Gap Constraints) Consider a sequence $S = e_1 - e_2 - \dots - e_n$ and an occurrence $O = i_1, i_2, \dots, i_m$ of a subsequence S' , if $(i_{k+1} - i_k) \leq g + 1$, such that $|k| \in \{1, \dots, m - 1\}$, then S' for the occurrence O , fulfills the gap constraint of g . Moreover, fulfilling the gap constraint once, in a given sequence serves the condition of gap-constraint satisfaction. For example, if $g = 2$, then XY is a subsequence of XZY , but not $XZZZY$.

Now, let us consider $D = \{D_1, D_2, \dots, D_n\}$ as a set of sequences in a database, a sequential pattern P , and a gap-constraint of g , then the frequency of occurrences of P in D is given by $count_P(D, g)$, also known as the absolute frequency support of P in D . If there exists a frequency support threshold α and P satisfies a condition such as $count_P(D, g) \geq \alpha$, then P is said to be frequent in D , with a gap constraint of g .

Definition 3.3.1.4: (Gap constrained sequential contrast patterns) Given two sets of sequence datasets D^+ (positive sequences) and D^- (negative sequences), two thresholds α and

δ , and a maximum gap of g , a gap-constrained sequential contrast pattern P is required to satisfy the following conditions.

- (1) Positive Support: $count_P(D^+, g) \geq \alpha$
- (2) Negative Support: $count_P(D^-, g) \leq \delta$

Thus given D^+ , D^- , α , δ and g , mining the gap-constrained sequential patterns involves finding the set of all such subsequences that fulfill the given conditions from (1) to (2).

2) *Generation of Candidate Sequences*: Towards finding the set of all gap-constrained contrast sequential patterns, we employ the ConSGapMiner algorithm [18], which was earlier used to extract minimal distinguishing subsequences (MDS) with user-defined gap constraints. The method utilizes the depth first search (DFS) technique for the generation of candidate sequences. This is done by growing a lexicographic sequence tree (LST) as shown in the example in Figure 4. Each node in the LST embeds a subsequence, along with its positive and negative frequency supports. In addition, each node is a max-prefix of its children.

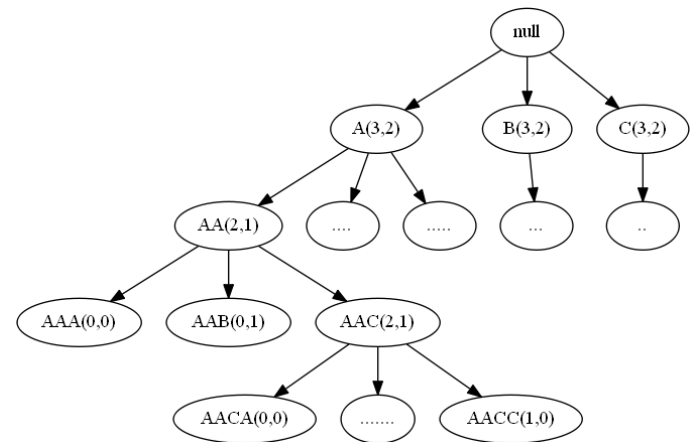


Fig. 4. A Lexicographic Sequence Tree (LST) growing candidate sequences using 3 symbols as A, B, C

Pruning non-minimal subsequences: After a sequence node is generated, if it satisfies the conditions (1) and (2), then the sequence node is not extended further. A supersequence of a potential contrast sequence is not minimal [18]. Thus, restricting the growth of sequences by a minimality condition, helps in the reduction of redundant patterns.

Pruning of infrequent subsequences: If a sequence node's positive frequency support is less than α (as specified in condition (1)), then the concerned node need not be extended. This is because, supersequences of an infrequent max-prefix are also infrequent.

3) *Gap Constraint Verification*: For the verification of gap-constraint satisfaction, we employed a bitmap representation reported earlier for checking gap-constraints [41]. The bitmap process is explained by an example, as shown in Table II. Let us consider verifying the gap constraint of XY in $XZXYZY$, given maximum gap g is set to 2. In the first step, all the occurrences of X in the concerned sequence are set to 1 (as shown in X_{index}). These are position indices given by 1 and 3. Later, $(g+1)$ index positions are set to 1 for each occurrences following X , separately as illustrated in rows 3 (given as 1_X)

and 4 (given as 2_X). Following this, the bit vectors in rows 3 and 4 go through a logical OR operation, as given in row 5. Subsequently, a logical AND operation is performed on the bit vectors in row 5 and for the occurrences of Y in row 6, to obtain a final bit vector, in row 7. An occurrence of 1 in the final bit vector (at row 7) indicates that the gap constraint of $g = 2$ was satisfied.

TABLE II
CHECKING GAP CONSTRAINT SATISFACTION OF XY IN XZXZY

	X	Z	X	Z	Y
Index	1	2	3	4	5
X_{index}	1	0	1	0	0
1_X	0	1	1	1	0
2_X	0	0	0	1	1
$1_X(OR)2_X$	0	1	1	1	1
Y	0	0	0	0	1
AND	0	0	0	0	1

Finally, a post-processing step is applied such that any supersequence of at least another shorter subsequence, is removed from the resulting set of contrast sequences. The algorithm for the generation of candidate sequences is provided by Algorithm I.

Algorithm 1 Generation of candidate sequences

$candGen(c, g, I, \delta, \alpha)$

```

1: Require :  $c$  – sequence,  $g$  – maximum gap,  $I$  –
  alphabet,  $\alpha$  – maximum positive support,  $\delta$  –
  minimum negative support
2:  $ds \leftarrow \phi$  { $ds$  holds the distinguishing children of  $c$ }
3: for  $i \in I$  do
4:   if  $c + i$  is not a supersequence of any sequence in  $ds$ 
     then
5:      $nc \leftarrow c + i$ 
6:      $supppos = SupportCount(nc, g, pos)$ 
7:      $suppneg = SupportCount(nc, g, neg)$ 
8:
9:     if  $supppos \geq \alpha AND suppneg \leq \delta$  then
10:       $ds \leftarrow ds \cup nc$ 
11:   else
12:     if  $supppos \geq \alpha$  then
13:        $candGen(c, g, I, \delta, \alpha)$ 
14:    $DS \leftarrow DS \cup ds$ 

```

IV. PREDICTION RESULTS

The sequential contrast pattern mining methodology was applied to both single-mode and multi-mode datasets, based on a clinical inclusion criteria, similar to principles used in [15]. From the MIMIC-II database, we extracted 253 segments (single mode) and 759 segments (multi-mode), which satisfied the criteria of hypotension. For the normotensive group, 274 segments were compiled for single mode whereas for multi-mode the exact number of segments varied from 133712 to 140006.

In addition, we also applied our techniques to the datasets provided by the Physionet 2009 AHE prediction challenge [5]. In particular, the AHE challenge datasets had also been

extracted from the MIMIC-II database in 2009. Our single-mode and multi-mode datasets tend to extend these datasets, since MIMIC-II has undergone multiple version updates, in the past 10 years. For the Physionet challenge, each of H1, H2, C1 and C2 groups consisted of 15 samples for training purposes. For test sets, Event I included 10 samples (H1=5, C1=5), while Event II had 40 (H=14, C=26). For the challenge data, an example training record like *a40439* contains a T_0 time-annotation, indicated as 18.30 on 04/09/2008 (T_0 was provided with each record). The time series data prior to T_0 is used for training purposes (treated as the observation window).

For the prediction of a record, a majority vote of contrast sequences is considered for the record to be treated as hypotensive. Single and multimode datasets extracted for the present study are available via <https://github.com/s-ghosh/hypotension>

A. Prediction performance on the two data sets

On the first data set, our 5-fold cross-validation classification results for both the single mode and multi-mode cases are summarized in Table III and IV. As can be noted, the classification results for the single mode executions are much better than multi-mode executions. This is because the single mode cross-validation accuracies are higher than multi-mode accuracies. A lower specificity in single mode executions can be attributed to the balanced nature of the single mode datasets. In contrast, the multi-mode datasets consist of a significantly higher percentage of instances, which are normotensive (for. e.g, 759 H to 140006 N). A sensitivity of 100% in our experiments, indicates that the sequential contrast method was able to predict all AHE instances correctly. Typically, the number of AHE instances are much fewer in comparison to non-AHE instances. As a result, the contrast pattern set generated due to the imbalance, can also consist of patterns which fulfill support conditions among non-AHE instances. Owing to this reason, contrast sets are highly capable of identifying positive instances. However, lower specificities reflect that a high percentage of false positives are also generated. Thus, our method demonstrates good performance when employed in the prediction of an AHE. This means sequential contrast patterns are effective in detecting hypotensive behaviour. However, since similar blood pressure patterns also exist across both population groups, a lot of negative instances are incorrectly classified as hypotensive. Similar experiments on MIMIC-II by Rocha et al [50] demonstrated a sensitivity of 82.8% and a specificity of 78.4%. In another study, Lee and Mark [15] also demonstrated highest accuracies of 76% for single-mode and 86% for multi-mode datasets extracted from MIMIC-II. Moreover, increasing the size of the observation window does not result in significant improvements in performance. Also, increasing gap intervals from 60 to 120 minutes lead to a drop in performance. Specifically, the hypotensive (positive) segments were always predicted correctly in both the modes.

Generally, retrospective EHR based population comparison studies tend to have imbalanced datasets, where the count of positive instances is very small as compared to the negative instances. As a possible enhancement, contrast pattern sets can be post-processed using multi-objective optimization methods

to obtain the most optimal combinations of contrast sequences for building models, which demonstrate better specificity, while reporting a higher classification performance.

For the Physionet 2009 challenge dataset, the test prediction results are presented in Table V. In Table VI, we provide a comparison of our results with the reported results from the Physionet 2009 challenge. As seen, models employing neural networks (GRNN, RPS-NN) and kernel methods like SVM are heavily dependent on several parameters, and can have performances over wide ranges [9], [10], [11]. Most of the other methods employed rules based on simple averaging measures and still performed fairly [8], [45]. Moreover, hidden markov models (HMM) for hypotension had reported a cross-validation accuracy close to 97% [49], which compares well with our cross-validation results too.

TABLE V
PHYSIONET 2009 AHE TEST PREDICTION CLASSIFICATION ACCURACIES FOR EVENTS I AND II GIVEN G=3

	Event I			Event II		
	S=3	S=4	S=5	S=3	S=4	S=5
L=8	5/10	7/10	7/10	23/40	23/40	32/40
L=9	5/10	7/10	9/10	23/40	25/40	33/40
L=10	5/10	7/10	10/10	25/40	32/40	36/40
L=11	5/10	7/10	10/10	25/40	32/40	36/40

TABLE VI
A COMPARISON OF CLASSIFICATION METHODS EMPLOYED FOR THE AHE PREDICTION PROBLEM. SEQUENTIAL PATTERNS REPORT COMPARABLE ACCURACIES AGAINST EXISTING METHODS

Method	Event I	Event II
GRNN	10/10	37/40
5-min average of diastolic ABP	10/10	37/40
MAP averaging Rule	10/10	36/40
5-min average of ABP	10/10	36/40
Linear Regression	10/10	36/40
Median of MAP	10/10	34/40
NN with feature selection	9/10	32/40
SVM	10/10	30/40
RPS-NN	2/10	25/40
Sequential Contrast Patterns	10/10	36/40

B. Discussion

A comparison of our results with the reported results from the Physionet 2009 challenge demonstrates our competitive classification performances against those models employing neural networks (GRNN, RPS-NN), kernel methods like SVM, hidden markov models and various other statistical measures [8],[9],[12]. Additionally, the effect of parameters like subsequence length (L), alphabet size (S) and maximum gap (G) are shown in Figure 5. As seen, the best performances were achieved using a maximum gap of 3, subsequence length of 10 and an alphabet of cardinality 5. A general trend is observed, where informative sequences could be extracted if the maximum gap constraint is iteratively increased. This has been demonstrated by Figure 5.

As seen, classification performances tend to improve with an increase in gap sizes. At the same time, a very large

gap size G, also means that two consecutive symbols for a sequential pattern have occurred over a wide range, where the size was G. Extremely large gap sizes can impede a proper interpretation of contiguous events in a sequence. Typically, patient events occur over a time span, covering multiple days. Thus, sequences may be clinically useful and unique, when considered for shorter time windows in the original patient timeline, with multiple days. For larger cohorts, finding out an optimal gap is dependent on the resolution of the time series (i.e. the sampling frequency). Typically, for detecting differential blood pressure patterns, effective gap sizes can be decided based on their ability to capture clinically meaningful and informative episodes, spanning over shorter windows. In addition, increasing S provides more number of discrete cut points for MAP, and enables the algorithm to capture patterns which characterize more fluctuations in the blood pressure. Thus, for cases with S=5, the algorithm is able to find a more expressive pattern, than for S=3. Hence, selecting an alphabet size of 5 turned out to be an optimal choice, both in terms of the discretization of blood pressure range as well as keeping the algorithmic running costs within limits. This also contributes to making improved predictions. Thus, finding interesting sequences is highly dependent on the use of various parameters like the number of symbols, length of subsequence and gap sizes. Generally, the selection of appropriate parameter values like L (pattern length), G (gap size) and S (alphabet) tends to affect the cardinality of the set of discovered patterns and the algorithmic running time. Thus, extracting minimally expressive shorter sequences allows the algorithm to restrict the running time as well as identify patterns, which are clinically important and appear in longer sequences.

In contrast to our method, the 5 minute averaging measures are statistical features obtained from a 5 minutes window prior to the immediate occurrence of an AHE. Thus, a major difference lies in the fact that our method considers a wider window of 30 and 60 minutes, prior to the onset of AHE [8]. This also indicates that a method, which is effective in performing predictions using wider time windows may be more suitable in a real time scenario, in comparison to statistical measures obtained from a 5 minutes window (prior to AHE). In this context, better results from the 5 minutes timespan prior to an AHE, may be due to temporal proximity to the onset of an AHE. For methods employing neural networks, both GRNN and RPS-NN report 10/10, 2/10 (for Event 1) and 37/40, 25/40 (for Event II). These methods tend to be strongly dependent on the tuning of parameter, as was also discussed by the authors [9]. The contrast mining method, on the other hand, helps to extract discretized sequential representations of the MAP time series, which provide the maximum support towards the occurrence of an AHE. These patterns are later useful, to not only predict an AHE for an unknown record, but for further clinical interpretation by domain experts. Our results indicate that sequential contrast patterns are capable of extracting informative symbolic episodes, which may be employed for both AHE risk prediction and understanding of hemodynamic behaviour towards effective analyses of sequential episodes, that may be indicative of medical symptoms.

TABLE III
SINGLE MODE CLASSIFICATION PERFORMANCE WITH 10 SYMBOLS

	Gap Interval = 60 minutes		Gap Interval = 120 minutes	
	ObWin = 0.5 h	ObWin = 1 h	ObWin = 0.5 h	ObWin = 1 h
Sensitivity	100%	100%	100%	100%
Specificity	65.85%	68.29%	61.44%	62.19%
Accuracy	82.27%	83.54%	79.87%	80.37%

TABLE IV
MULTI MODE CLASSIFICATION PERFORMANCE WITH 15 SYMBOLS

	Gap Interval = 60 minutes		Gap Interval = 120 minutes	
	ObWin = 0.5 h	ObWin = 1 h	ObWin = 0.5 h	ObWin = 1 h
Sensitivity	100%	100%	100%	100%
Specificity	81.19%	80.76%	79.36%	74.79%
Accuracy	81.30%	80.88%	79.48%	74.94%




Fig. 5. Effect of parameters L and G on the performance (A) For Event I, (B) For Event II

V. EXAMPLES AND CLINICAL SIGNIFICANCE OF SEQUENTIAL CONTRAST PATTERNS

Acute hypotension is one of the most dangerous clinical conditions that frequently occurs in an ICU and can cause serious renal, cerebral and myocardial hypoxic damage. Existing medical interventions are reactive (i.e after an AHE has been triggered), for recommending treatment of underlying causes. In contrast, early bedside detection of AHEs can enable the development of life-saving interventions. Clinical interventions to treat AHE attempt to restore the physiological status of the body by targeting recommended BP values, increasing fluid and salt intakes, administration of vasoactive agents and so on [39], [40]. The AHE definition considered in the current study, utilizes hypotension thresholds reported in previous studies [15],[5]. Although ranges between 65-75 mmHg have also been reportedly used for defining hypotension, definitions for AHE time periods may also vary from 1 to 60 minutes, depending on the objective of the study. However, drops in blood pressure within smaller time spans (as indicated by monitoring systems), may not always indicate an AHE. Such changes may be due to monitoring errors or physiological changes caused by normal human activity. Hence, a larger time window of 30 minutes is a suitable definition for capturing AHE related information. Taking forward the suggested inclusion criteria

for an AHE, we additionally employed the widely used ICD-9 code of hypotension to extract clinical records from MIMIC-II. The ICD-9 coding system describes a disease classification scheme used to monitor population group health situations for general epidemiological, health management purposes and clinical usage. The extracted datasets were sourced from the MIMIC-II repository, which tends to provide further credence to the study.

A. Sequential pattern examples

Our sequential pattern mining algorithm can discover simple-to-understand clinical symbolic subsequences. These subsequences can be treated as evidence while diagnosing for diseases. Even though methods such as neural networks and SVM demonstrate competitive prediction performances, they heavily dependent on non-linear kernel functions and parameters. But, our sequential pattern mining methods extract signatures of clinical episodes in the form of symbolic patterns.

In this study, we were able to mine a set of discretized sequential patterns like ABAEDBBBCA, which were prominent in acute hypotensive patients. Examples of representative sequential blood pressure patterns for hypotension are as reported in Table VII. For example, in the case of ABAEDBB-

TABLE VII
REPRESENTATIVE EXAMPLES OF EXTRACTED AHE SEQUENTIAL PATTERNS

..D..E..D..E..D..A..B..C..D..C..
..D..C..E..D..C..B..C..D..C..D..
..B..C..D..C..A..C..D..E..D..C..
..D..C..E..C..A..C..D..E..D..E..
..C..A..B..A..E..E..C..B..C..D..
..A..B..A..E..D..B..B..B..C..A..
..E..C..B..A..B..A..B..C..D..
..A..B..B..D..E..C..B..C..D..

BCA, the sequence indicates that the mean arterial pressure follows the given pattern trajectory among a majority of AHE patients. The given symbols indicate that the mean arterial pressure time series region was divided into 5 equiprobable regions (given by A, B, C, D, E) from 0 to 200 mmHg. The example pattern illustrates that the blood pressure time series followed a situation where majority of the AHE patients record an episode of events represented by the MAP value in a particular sequential order of blood pressure regimes demonstrated symbolically as follows - $A \leq B \leq A \leq E \leq D \leq B \leq B \leq B \leq B \leq C \leq A$. Thus, each sequential pattern describes a train of clinical events, represented by the specific blood pressure regimes, categorised by discrete symbols.

B. Pattern visualization and clinical interpretation

Interpretive sequential representations can be extremely useful to clinicians for understanding the sequence of physiological states that a patient passes through, before developing a critical condition. Such interpretations can help establish potential combinations of observable physiological sequences, that precede AHE. Generally, the objective of clinical studies involves the estimation of causal relationships between selected clinical variables and disease specific laboratory test outcomes. Given temporal data for clinical variables, sequential patterns of specific clinical variables can aid in the interpretation of complex relationships between variables and patient specific outcomes. Towards this objective, general visual trends may be inferred from gap constrained sequences as shown in Figure 6. Thus, sequential patterns can have immense potential in the exploration of underlying clinical relationships to facilitate personalized treatments. Accordingly, similar studies have also claimed that the visual exploration of sequential and temporal patterns in clinical patient data can significantly aid in clinical decision making [53].

Moreover, mining of complex contrast sequences in hypotensive patient groups can aid in the development of interesting clinical hypotheses such as the detection of a succession of clinical events prior to the onset of AHEs. Thus, extracting sequential contrast patterns can guide clinical decision-making towards the effective investigation of hypotensive events. In addition, the proposed methodology is flexible enough to also accommodate clinician-defined constraints.

VI. CONCLUSION

The current study investigated the application of a novel sequential contrast pattern mining methodology for predicting acute hypotensive episodes in an ICU. Our study demonstrates that research on the mining of informative sequential patterns can be of significant clinical value to concerned stakeholder in a clinical setting. In addition to demonstrating the classification performance, we also established the existence of gap-constrained symbolic subsequences, which have strong clinical interest to practitioners. Since the data encoded of a patient's journey is inherently temporal in nature, sequences have the capability to uncover numerous hidden patterns, which are otherwise not visible. As part of a knowledge discovery process, the contrast pattern mining method extracts patterns, which collectively help in the prediction of an AHE. A real-time application of the reported strategy can help derive significant sequential patterns of interest, which could be translated into a complex sequence of clinical events. A higher frequency of the occurrence of complex contrast sequences while comparing hypotensive and normotensive patient groups may be beneficial to a clinician to develop a clinical hypothesis relating to a succession of clinical events leading to an AHE. Extracting sequential patterns from hypotensive patient groups can inform medical decision-making towards the diagnosis and investigation of AHEs. Thus, significant patterns are a potential source for launching further data driven investigations validated by randomized clinical trials. Such patterns can also be employed in conjunction with multiple types of clinical features for the construction of accurate AHE prediction systems. In summary, the sequential contrast pattern mining approach described in this work well relates to the expectations of evidence-based medicine.

ACKNOWLEDGMENT

This work was partly supported by Australia Research Council Discovery Project (DP130102124) and the R01 grant EB001659 from the National Institute of Biomedical Imaging and Bioengineering (NIBIB) of the National Institutes of Health NIH. Mengling Feng is supported by the Agency for Science, Technology and Research Graduate Scholarship, Singapore. Shameek Ghosh is supported by the Dr A.P.J. Abdul Kalam & UTS Doctoral Scholarship, Australia.

REFERENCES

- [1] J. Boldt, Clinical review: Hemodynamic monitoring in the intensive care unit, *Critical Care*, vol 6, number 1, pp. 52, 2002
- [2] M. R. Pinsky, Hemodynamic evaluation and monitoring in the ICU, *CHEST Journal*, 132(6), 2020-2029. 2007
- [3] M. Saeed, M. Villarroel, A.T. Reinsner, G. Clifford, L.W. Lehman, G. Moody et al., Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): a public-access intensive care unit database, *Critical care medicine*, 39(5), 952, 2011.
- [4] H. Calkins and D.P. Zipes, Hypotension and syncope, *Heart Disease. A Textbook of Cardiovascular Medicine*. Philadelphia, PA: Saunders, 909-19, 2005.
- [5] G. B. Moody and L. H Lehman, Predicting acute hypotensive episodes: The 10th annual physioNet/computers in cardiology challenge, In *Computers in Cardiology*, (pp. 541-544), 2009.
- [6] M. Saeed, and R. Mark, A novel method for the efficient retrieval of similar multiparameter physiologic time series using wavelet-based symbolic representations. In *AMIA Annual Symposium Proceedings*, Vol. 2006, p. 679, American Medical Informatics Association, 2006

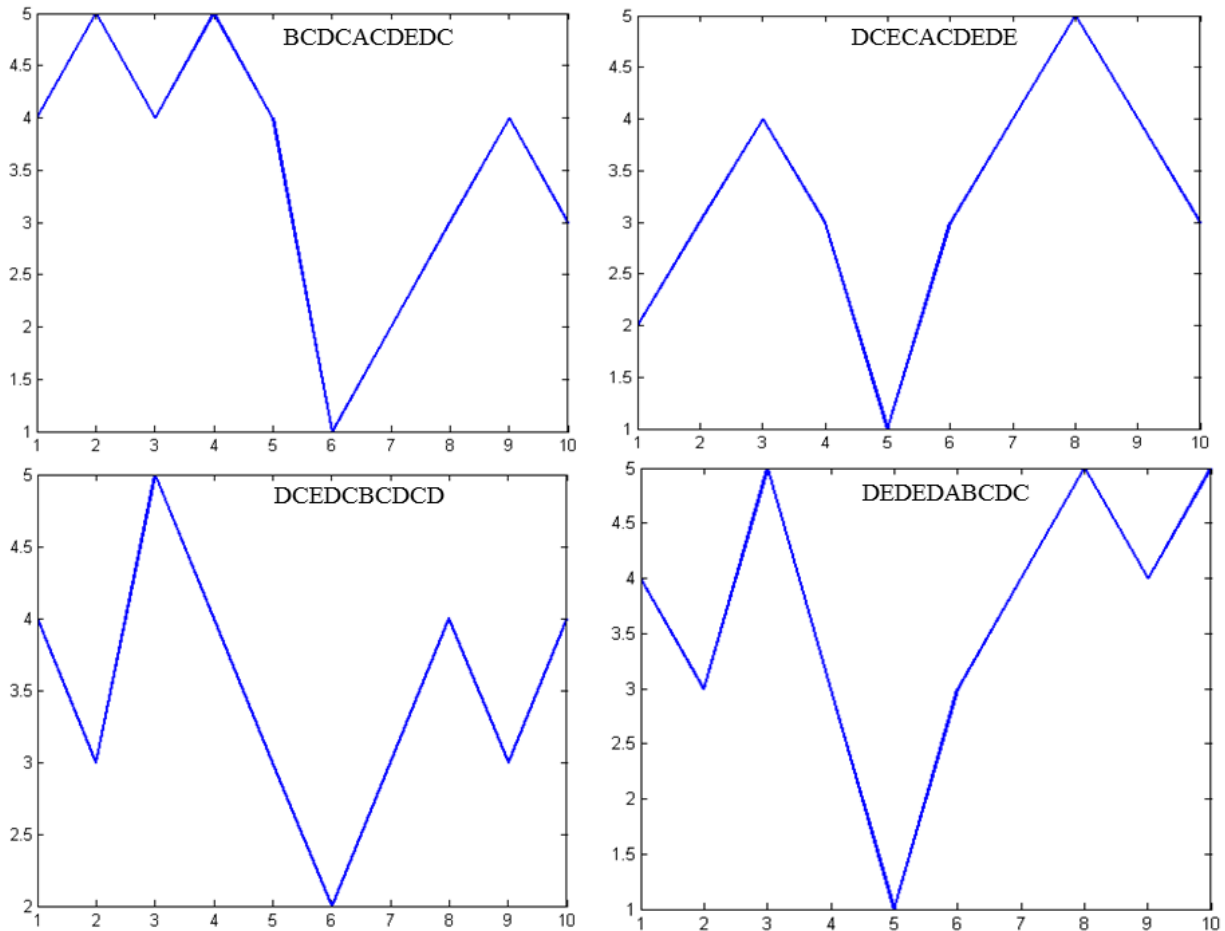


Fig. 6. Inferring Visual Trends from Sequential Contrast Patterns Examples for AHE (A=1, B=2, C=3, D=4, E=5)

- [7] A. Ghaffari, M.R. Homaeinezhad, M. Atarod, M. Akraminia, Parallel processing of ECG and blood pressure waveforms for detection of acute hypotensive episodes: a simulation study using a risk scoring model, *Computer methods in biomechanics and biomedical engineering*, 13(2), 197-213, 2010.
- [8] X. Chen, D. Xu, G. Zhang, R. Mukkamala, Forecasting acute hypotensive episodes in intensive care patients based on a peripheral arterial blood pressure waveform, *In Computers in Cardiology*, pp. 545-548, 2009.
- [9] J. Henriques, T. R. Rocha, Prediction of acute hypotensive episodes using neural network multi-models, *In Computers in Cardiology*, (pp. 549-552), 2009.
- [10] M. A. Mneimneh, R. J. Povinelli, A rule-based approach for the prediction of acute hypotensive episodes, *In Computers in Cardiology*, (pp. 557-560), 2009.
- [11] F. Jousset, M. Lemay, J. M. Vesin, Computers in cardiology/physioNet challenge 2009: Predicting acute hypotensive episodes, *Computers in Cardiology*, 36, 637-640, 2009.
- [12] P. Langley, S. T. King, D. Zheng, E. J. Bowers, K. Wang, J. Allen, A. Murray, Predicting acute hypotensive episodes from mean arterial pressure. *In Computers in Cardiology*, (pp. 553-556), 2009.
- [13] F. Wang, N. Lee, J. Hu, J. Sun, S. Ebadollahi, A. F. Laine, A Framework for Mining Signatures from Event Sequences and Its Applications in Healthcare Data, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(2), 272-285, 2013.
- [14] Z. Syed, C. Stultz, M. Kellis, P. Indyk, J. Guttag, Motif discovery in physiological datasets: a methodology for inferring predictive elements, *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 4(1), 2, 2010.
- [15] J. Lee, R. G. Mark, An investigation of patterns in hemodynamic data indicative of impending hypotension in intensive care, *Biomedical engineering online*, 9(1), 62, 2010.
- [16] J. Lin, E. Keogh, S. Lonardi, B. Chiu, A symbolic representation of time series, with implications for streaming algorithms, *In Proceedings of the 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery*, (pp. 2-11), 2003.
- [17] G. Dong, J. Li, Efficient mining of emerging patterns: Discovering trends and differences, *In Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining*, (pp. 43-52), 1999.
- [18] X. Ji, J. Bailey, G. Dong, Mining minimal distinguishing subsequence patterns with gap constraints, *Knowledge and Information Systems*, 11(3), 259-286, 2007.
- [19] American Medical Association, International classification of diseases, 9th revision, clinical modification: physician ICD-9-CM, 2005: volumes 1 and 2, color-coded, illustrated (Vol. 1), 2004.
- [20] A. Ghaffari, M. R. Homaeinezhad, M. Atarod, M. Akraminia, A methodology for prediction of acute hypotensive episodes in ICU via a risk scoring model including analysis of ST-Segment Variations, *Cardiovascular Engineering*, 10(1), 12-29, 2010.
- [21] L. Mayaud, P. S. Lai, G. D. Clifford, L. Tarassenko, L. A. Celi, D. Annane, D. Dynamic data during hypotensive episode improves mortality predictions among patients with sepsis and hypotension, *Critical care medicine*, 41(4), 954, 2013.
- [22] F. Giza, J. Van Eyck, G. Meyfroidt, Predictive data mining on monitoring data from the intensive care unit, *Journal of clinical monitoring and computing*, 27(4), 449-453, 2013.
- [23] P. A. Low, Prevalence of orthostatic hypotension, *Clinical Autonomic Research*, 18(1), 8-13, 2008.
- [24] F. Schmid, M. S. Goepfert, D. A. Reuter, Patient monitoring alarms in the ICU and in the operating room, *Crit Care*, 17(2), 216, 2013.
- [25] S. Nizami, J. R. Green, C. McGregor, Implementation of artifact detection in critical care: a methodological review, *IEEE Reviews in Biomedical Engineering*, 6, 127-142, 2013.
- [26] L. J. Eshelman, K. P. Lee, J. J. Frassica, W. Zong, L. Nielsen, M. Saeed, Development and evaluation of predictive alerts for hemodynamic

- instability in ICU patients, In AMIA Annual Symposium Proceedings (Vol. 2008, p. 379). American Medical Informatics Association, 2008.
- [27] R. Donald, T. Howells, I. Piper, I. Chambers, G. Citerio, P. Enblad et al., Early warning of EUSIG-defined hypotensive events using a Bayesian Artificial Neural Network, *Acta Neurochirurgica Supplementum*, Volume 114, 2012, pp 39-44, 2012.
- [28] P. E. Owens, S. P. Lyons, E. T. O'Brien, Arterial hypotension: prevalence of low blood pressure in the general population using ambulatory blood pressure monitoring, *Journal of human hypertension*, 14(4), 243-247, 2000.
- [29] L. A. Celi, L. C. Hinske, G. Alterovitz, P. Szolovits, An artificial intelligence tool to predict fluid requirement in the intensive care unit: a proof-of-concept study, *Crit Care*, 12(6), R151, 2008.
- [30] G. Meyfroidt, F. Giza, D. Cotte, W. De Becker, K. Van Loon, J. M. Aerts et al., Computerized prediction of intensive care unit discharge after cardiac surgery: development and validation of a Gaussian processes model, *BMC medical informatics and decision making*, 11(1), 64, 2011.
- [31] R. Moskovitch, C. Walsh, G. Hripsack, N. Tatonetti, Prediction of Biomedical Events via Time Intervals Mining, *ACM KDD Workshop on Connected Health in Big Data Era*, 2014.
- [32] D. Patnaik, P. Butler, N. Ramakrishnan, L. Parida, B. J. Keller, D. A. Hanauer, Experiences with mining temporal event sequences from electronic medical records: initial successes and some challenges, In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, (pp. 360-368), 2011.
- [33] I. Batal, D. Fradkin, J. Harrison, F. Moerchen, M. Hauskrecht, Mining recent temporal patterns for event detection in multivariate time series data, In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, (pp. 280-288), 2012.
- [34] T. Toma, A. Abu-Hanna, R. J. Bosman, Discovery and integration of organ-failure episodes in mortality prediction, In *Artificial Intelligence in Medicine* (pp. 97-106), Springer Berlin Heidelberg, 2007.
- [35] L. Sacchi, R. Bellazzi, C. Larizza, R. Porreca, P. Magni, Learning rules with complex temporal patterns in biomedical domains, In *Artificial Intelligence in Medicine* (pp. 23-32). Springer Berlin Heidelberg, 2005.
- [36] R. Bellazzi, C. Larizza, P. Magni, R. Bellazzi, Temporal data mining for the quality assessment of hemodialysis services, *Artificial intelligence in medicine*, 34(1), 25-39, 2005.
- [37] R. Bellazzi, F. Ferrazzi, L. Sacchi, Predictive data mining in clinical medicine: a focus on selected methods and applications, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(5), 416-430, 2011.
- [38] R. Moskovitch, Y. Shahar, Medical temporal-knowledge discovery via temporal abstraction, In AMIA Annual Symposium Proceedings (Vol. 2009, p. 452). American Medical Informatics Association, 2009.
- [39] C. Shibao, L. A. Lipsitz, I. Biaggioni, ASH position paper: evaluation and treatment of orthostatic hypotension, *The Journal of Clinical Hypertension*, 15(3), 147-153, 2013.
- [40] J. Takala, Should we target blood pressure in sepsis?, *Critical care medicine*, 38(10), S613-S619, 2010.
- [41] J. Ayres, J. Flannick, J. Gehrke, T. Yiu, Sequential pattern mining using a bitmap representation, In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 429-435), 2002.
- [42] S. Chan, B. Kao, C. L. Yip, M. Tang, Mining emerging substrings. In *Eighth International Conference on Database Systems for Advanced Applications*, (pp. 119-126), 2003.
- [43] Z. Xing, J. Pei, E. Keogh, A brief survey on sequence classification, *ACM SIGKDD Explorations Newsletter*, 12(1), 40-48, 2010.
- [44] N. H. Shah, J. D. Tenenbaum, The coming age of data-driven medicine: translational bioinformatics' next frontier, *Journal of the American Medical Informatics Association*, 19(e1), e2-e4, 2012.
- [45] P. A. Fournier, J. F. Roy, Acute hypotension episode prediction using information divergence for feature selection, and non-parametric methods for classification, In *Computers in Cardiology*, 2009 (pp. 625-628), 2009.
- [46] R. J. Anderson, Plumbing the depths of blood pressure: Hypotensive hemorrhage and acute kidney injury*, *Critical care medicine*, 39(9), 2196-2197, 2011.
- [47] D. C. Angus, T. van der Poll, Severe sepsis and septic shock, *New England Journal of Medicine*, 369(9), 840-851, 2013.
- [48] H. H. Awad, F. A. Anderson Jr, J. M. Gore, S. G. Goodman, R. J. Goldberg, Cardiogenic shock complicating acute coronary syndromes: insights from the Global Registry of Acute Coronary Events, *American heart journal*, 163(6), 963-971, 2012.
- [49] A. Singh, T. Tamminedi, G. Yosiphon, A. Ganguli, J. Yadegar, Hidden Markov Models for modeling blood pressure data to predict acute hypotension, In *IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*, (pp. 550-553), 2010.
- [50] T. Rocha, S. Paredes, P. De Carvalho, J. Henriques, Prediction of acute hypotensive episodes by means of neural network multi-models, *Computers in biology and medicine*, 41(10), 881-890, 2011.
- [51] M. Ghassemi, L. H. Lehman, J. Snoek, S. Nemati, Global Optimization Approaches for Parameter Tuning in Biomedical Signal Processing: A Focus of Multi-scale Entropy, In *Computing in Cardiology Conference (CinC)*, pp.993-996, 2014.
- [52] A. Perer, F. Wang, Frequency: Interactive mining and visualization of temporal frequent event sequences, In *Proceedings of the 19th international conference on Intelligent User Interfaces* (pp. 153-162). 2014.
- [53] D. Gotz, F. Wang, A. Perer, A methodology for interactive mining and visual analysis of clinical event patterns using electronic health record data, *Journal of biomedical informatics*, 48, 148-159, 2014.
- [54] V. S. Tseng, C. H. Lee, Effective temporal data classification by integrating sequential pattern mining and probabilistic induction, *Expert Systems with Applications*, 36(5), 9524-9532, 2009.
- [55] V. S. Tseng, C. H. Lee, Cbs: A new classification method by using sequential patterns, In: *SDM 2005: Proc. of the 2005 SIAM International Data Mining Conference*, 2005.