Computers in Cardiology / Physionet Challenge 2009: Predicting Acute Hypotensive Episodes

F Jousset, M Lemay, JM Vesin

Laboratoire de Traitement des Signaux 1, EPFL, Switzerland

Abstract

The goal of the Computers in Cardiology / Physionet Challenge 2009 is to predict which patients will experience acute hypotensive episode within a forecast window of one hour. In our study, statistically robust features extracted from the supplied training set were defined. A Support Vector Machine was used to classify these features. In this paper, we present our method, results and conclusion about this statistical approach.

1. Introduction

Acute hypotensive episode (AHE), in which a patient's arterial blood pressure decreases to an abnormally low level, may result in dangerous complications and even death [1]. Predicting the onset of AHE would be of clear benefit to a patient's outcome and reduce his stay in intensive care unit (ICU).

Predicting AHE is an ongoing topic and is currently explored by Avert-IT. The european Avert-IT project proposes the development of a system architecture that can automatically predict adverse hypotensive events over a useful timescale, without the intervention of a health-care professional. The prediction task is performed through a Novel Bayesian Neural Network [2].

The development of methods to identify patients at risk of AHE is the topic of the Computers in Cardiology / Physionet Challenge 2009 [3]. The Challenge is composed of two events. The first event focuses on distinguishing between two groups of ICU patients who were currently under pressor medication: patients who experience an AHE (H1), and patients who do not (C1). These two groups represent extremes of AHE-associated risk. The second event aims at addressing the broad question of predicting AHE within the forecast window in a population in which about a third of the patients experience AHE (H2 and C2). The forecast window is defined as the one-hour period immediately following a specified time T0. In the test sets, the forecast window (and indeed all data following T0) are withheld, and the forecast must be made using only infor-
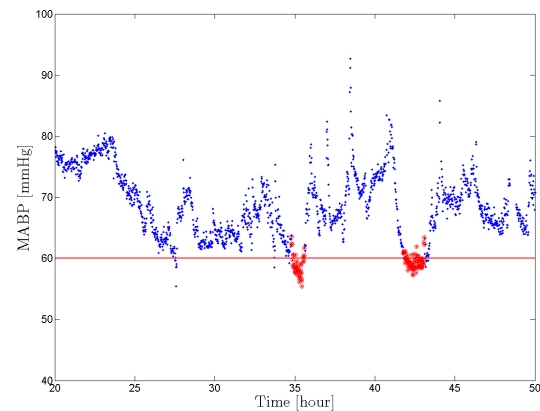

Figure 1 – Mean arterial blood pressure (H1 group). Two AHE below the event threshold at 60 mmHg are shown in red circles

mation available before T0.

The Challenge dataset was chosen from the MIMIC II database [4]. It included, for each case, a time series of mean arterial blood pressure (MABP) at one-minute intervals. Each sample of the series is an average of the blood pressure measured in the radial artery over the previous minute. Given such a time series, an AHE was defined for the purposes of the Challenge as any period of 30 minutes or more during which at least 90% of the MABP measurements were at or below 60 mmHg. Figure 1 shows a representative exemple of MABP in which two AHE occured.

We approached the Challenge from a statistical point of view and used a support vector machine (SVM) for the classification.

2. Methods

The proposed approach takes into consideration two aspects of the AHE detection task, namely the parameter selection and the classification block.
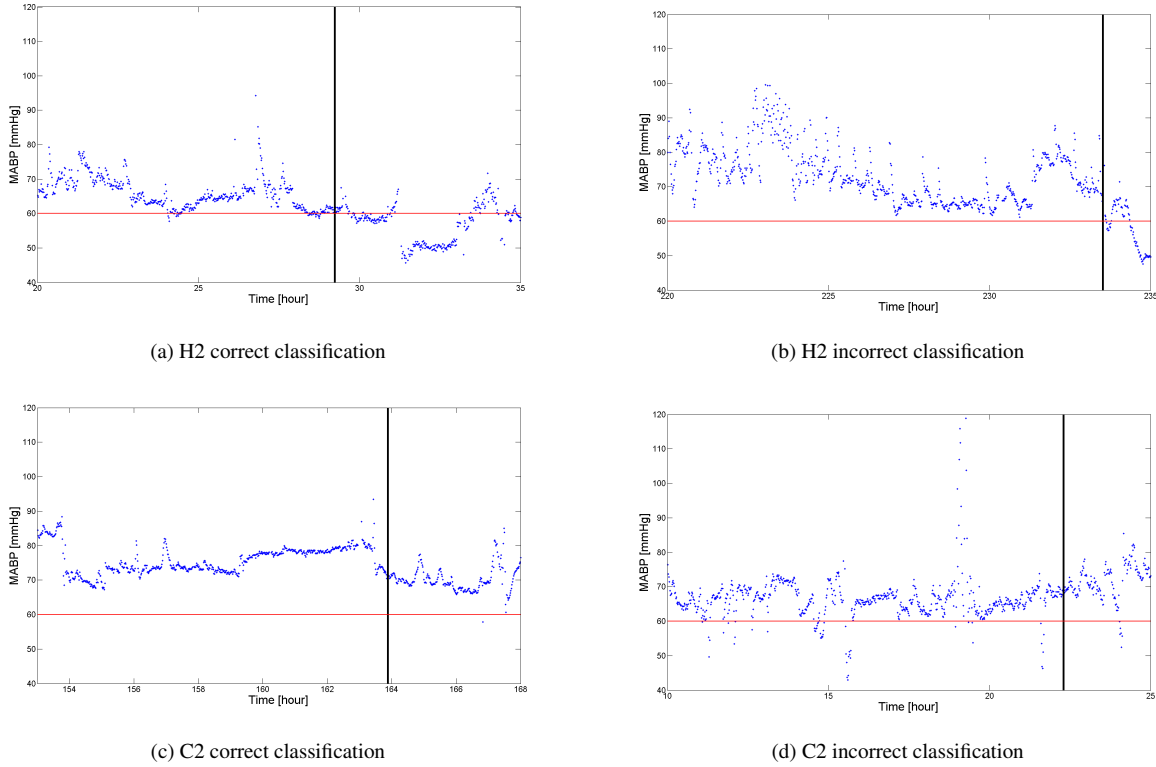
(a) H2 correct classification

(b) H2 incorrect classification

(c) C2 correct classification

(d) C2 incorrect classification

Figure 2 – Representatives examples of the MABP from the two training groups H2 and C2, the vertical line represents the time T0. The two subfigures on the left hand side show records from each group correctly classified as the two other show records incorrectly classified.

2.1.   Parameter selection

2.1.1.   Observations

A first visual inspection of the training sets gave us some informations about the two groups. The groups experiencing AHE after T0 seemed to have a MABP closer from the threshold of 60 mmHg and to be less variable on short windows preceding T0. Figure 2(a) and (c) shows these trends in records from each group. The horizontal line defines the AHE threshold as the vertical line defines the beginning of the forecast window T0. The two right hand side subfigures (Figure 2(b) and (d)) show counter-examples. In figure 2(b), the MABP suddenly dropped just after T0. In figure 2(d) the MABP crossed several time the threshold while exhibiting a great variability, but never remained long enough under the threshold to reach the AHE criteria.

2.1.2.   Features

We focused on the signals always available in the training and the test sets: the heart rate, the respiration, the diastolic, systolic and mean arterial blood pressure. Most records included a variety of additional vital signs signals, like respiration rate and saturation of peripheral oxygen (S_pO_2). Since those additional signals were not available for every record, we let them apart. The numerical time series of vital signs were sampled once per minute in the training sets and once per second in the test sets.

The systolic arterial blood pressure (SABP) is the maximum pressure when the heart contracts and blood begins to flow. The diastolic arterial blood pressure (DABP) is the minimum pressure occurring between heartbeats. The mean arterial blood pressure (MABP) is a combination of the two above quantities, most often calculated as:

$$DABP + \frac{SABP - DABP}{3} \qquad (1)$$

From our global observations, we extracted from the available signals statistical parameters, such as the mean, the standard deviation, the skewness and the kurtosis. We also computed robust statistics (median and median absolute deviation) in order to be less sensitive to outliers. The slope of the signals was computed using robust regression [6]. Those features were computed on signals of various lengths preceding the forecast window. The occurrence of

an episode of hypotension preceding the forecast window was marked in as an additional parameter.

2.1.3. Sub-selection

Using a leave-one-out cross validation (LOOCV) on the training sets, the parameters achieving the best classification rate were kept [5]. A LOOCV strategy is a common choice when the training sets are small (Table 1).

The features selected for the first event were the median MABP and the median absolute deviation (mad) MABP over the last 2 hours. The selection of features for the second event also included the median and the mad of the MABP as well as the slope of the MABP over the last 30 minutes.

	Event 1		Event 2	
	H1	C1	H2	C2
Training	15	15	15	15
Test	5	5	[10,16]	[24,30]

Table 1 – Number of record per challenge dataset

2.2. Acute hypotension episodes prediction

A support vector machine (SVM) approach was used to classify the features and predict the occurrence of an AHE [7]. The SVM separates the data by an optimal hyperplane. Optimization is obtained by maximizing the margin, i.e. the distance between the hyperplane and the nearest data points of each class, called support vectors. We used a SVM with a linear kernel for both events.

For the event 1, using the median and the mad of the MABP over a 2-hour period preceding T0 resulted in a correct classification of 29 training records out of 30. Figure 3 shows an example of the hyperplane clearly separating the two training groups H1 (AHE group) and C1. The support vectors are circled in black.

For the event 2, the addition of the slope of the MABP over the last 30 minutes, increased correct classification from 20 to 22 out of 30. The addition of other features did not improve the classification performance.

3. Results

The Computers in Cardiology Challenge scored the various methods, based on the percentage of correct classification on the test set for both events. For the first event, a total of 10 records and for the second one, a total of 40 records were provided (Table 1). To address the classification, the test records were resampled to one sample per minute, as the training sets were.
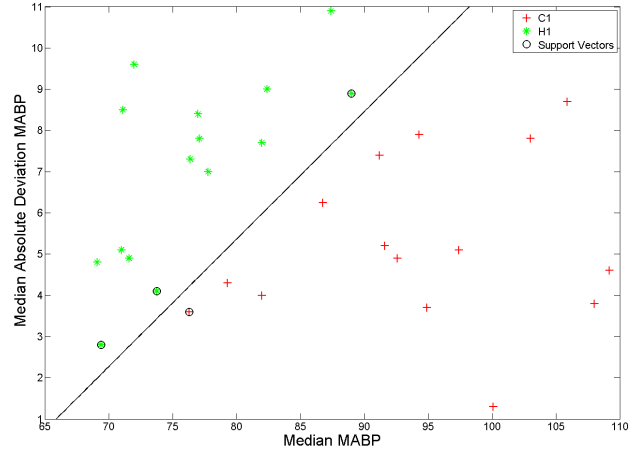

Figure 3 – Classification of the training set A, based on the median absolute deviation and the median of the MABP.

H1 (AHE) record #	C1 (no AHE) record #
101	103
102	105
104	106
109	107
110	108

Table 2 – Event 1, classification test set A

The classification was performed by the trained SVM structure and used only the information available before the forecast window. On the test set A, corresponding to the event 1, we correctly classified the 10 records. Exactly five of them were associated to the group H1 (i.e. will experience AHE in the forecast window). The record classification is detailed in Table 2.

In the second event, the exact number of records per group was not known. The number of records in group H2 could range from 10 to 16. We classified 16 records from the test set B in the group H2 and obtained a correct classification of 75% (Table 3).

		Actual Class	
		C2	H2
Predicted	C2	20	4
Class	H2	6	10

Table 3 – Event 2, confusion matrix

4. Discussion and conclusions

Using only information from the MABP, we performed well on the first event, where the two groups were well separated. Our method yielded a limited performance on the second event, that involves patients with a wider variety of histories and thus larger variability in the data.

In both events, robust statistics proved to perform better than standard ones. Furthermore, none of the features from other vital signs improved our classification rate.

Our analysis was only based on the numerical signals. A more detailed study could also investigate any change in the waveform signal prior to an AHE. Further studies may also take into account the medication delivered and medical history of the patients, in order to possibly improve in AHE prediction.

Acknowledgements

References

[1] Manley G, Knudson M, Morabito D, Damron S, Erickson V, Pitts L. Hypotension, Hypoxia, and head injury. ARCH SURG 2001 136:1118–1123.

[2] Stell A, Sinnott R, Jiang J, Donald R, Chambers I, Citerio G, Enblad P, Gregson B, Howells T, Kiening K, Nilsson P, Ragauskas A, Sahuquillo J, Piper I. Federating distributed clinical data for the prediction of adverse hypotensive events. Phil Trans R Soc A 2009 367:2679–2690.

[3] Moody GB, Lehman LH. Predicting acute hypotensive episodes: the 10th annual PhysioNet/Computers in Cardiology Challenge. Computers in Cardiology 2009;36.

[4] The MIMIC II Project database via the Physionet website. http://www.physionet.org/physiobank/database/mimic2db/

[5] Fukunaga K, Hummels DM. Leave-One-Out Procedures for Nonparametric Error Estimates. IEEE Trans. Pattern Anal. and Machine Intell. 1989 4(11):421–423

[6] Holland P. W., R. E. Welsch. Robust Regression Using Iteratively Reweighted Least-Squares. Communications in Statistics: Theory and Methods, A6, 1977, pp. 813-827.

[7] Cristianini N, and Shawe-Taylor J. Support Vector Machines and other kernel-based learning methods. Cambridge University Press, 2000.

Address for correspondence:

Florian Jousset
EPFL-STI-IEL-LTS1
ELD 231 Station 11
1015 Lausanne, Switzerland
florian.jousset@epfl.ch