When Less is Better: A Summarization Technique that Enhances Clinical Effectiveness of Data

P Durga*
Amrita Vishwa Vidyapeetham

Rahul Krishnan Pathinarupothi†
Amrita Vishwa Vidyapeetham

Ekanath Srihari Rangan‡
Amrita Vishwa Vidyapeetham

Prakash Ishwar§
Boston University

ABSTRACT

The increasing number of wearable sensors for monitoring of various vital parameters such as blood pressure (BP), blood glucose, or heart rate (HR), has opened up an unprecedented opportunity for personalized real-time monitoring and prediction of critical health conditions of patients. This, however, also poses the dual challenges of identifying clinically relevant information from vast volumes of sensor time-series data and of storing and communicating it to health-care providers especially in the context of rural areas of developing regions where communication bandwidth may be limited. One approach to address these challenges is data summarization, but the danger of losing clinically useful information makes it less appealing to medical practitioners. To overcome this, we develop a data summarization technique called RASPRO (Rapid Active Summarization for effective PROgnosis), which transforms raw sensor time-series data into a series of low bandwidth, medically interpretable symbols, called "motifs", which measure criticality and preserve clinical effectiveness benefits for patients. We evaluate the predictive power and bandwidth requirements of RASPRO on more than 16,000 minutes of patient monitoring data from a widely used open source challenge dataset. We find that RASPRO motifs have much higher clinical efficacy and efficiency (20 – 90% improvement in F1 score over bandwidths ranging from 0.2–0.75 bits/unit-time) in predicting an acute hypotensive episode (AHE) compared to Symbolic Aggregate approXimation (SAX) which is a state-of-the-art data reduction and symbolic representation method. Furthermore, the RASPRO motifs typically perform as well or much better than the original raw data time-series, but with up to 15-fold reduction in transmission/storage bandwidth thereby suggesting that *less is better*.

*The first and second authors have contributed equally to this work.
†P. Durga and R.K. Pathinarupothi are with Amrita Center for Wireless Networks & Applications (AmritaWNA), Amrita School of Engineering, Amritapuri, Amrita Vishwa Vidyapeetham, India. email: rahulkrishnan@am.amrita.edu.
‡E. S. Rangan is with School of Medicine, Cochin, Amrita Vishwa Vidyapeetham, India.
§P. Ishwar is with the ECE department, Boston University, Boston, MA, USA 02215.

DH'18, April 23–26, 2018, Lyon, France

ACM ISBN 978-1-4503-6493-5/18/04…$15.00
https://doi.org/10.1145/3194658.3194674

CCS CONCEPTS

• **Information systems → Decision support systems**; *Mobile information processing systems*;

KEYWORDS

Predictive health monitoring, personalization, data analytics

ACM Reference Format:
P Durga, Rahul Krishnan Pathinarupothi, Ekanath Srihari Rangan, and Prakash Ishwar. 2018. When Less is Better: A Summarization Technique that Enhances Clinical Effectiveness of Data. In *DH'18: 2018 International Digital Health Conference, April 23–26, 2018, Lyon, France.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3194658.3194674

1 INTRODUCTION

We are witnessing an exponential growth in the adoption of vitals monitoring sensors among the general population as well as the patient community [4]. Despite offering multiple advantages to the patients and healthcare providers in terms of continuous monitoring and reduced healthcare delivery costs, the challenge lies in mining the large amount of data that comes from these sensors to make sensible and timely diagnosis of critical health conditions such as myocardial infarctions, strokes, hypotensive episodes, or syncopes. Towards this end, the research community has been actively looking at various data summarization techniques [11]. In addition to storage and communication bandwidth savings, data reduction can potentially improve prediction power by improving the signal-to-noise ratio and eliminating redundant and uninformative signal dimensions. However, a concern of medical practitioners is that data reduction may result in either loss of feature details that could be clinically relevant or may produce reduced representations that are difficult to interpret medically [11]. In this paper, we propose a novel technique, which by utilizing medical knowledge to summarize voluminous sensor data, not only retains medical interpretability and reduces bandwidth, but also improves the efficacy of predicting diseases.

This technique, that we call as RASPRO (Rapid Active Summarization for effective PROgnosis), transforms raw sensor data to patient health status "motifs". These motifs are easily readable health summaries of the patient, and they directly help doctors and caretakers gauge the urgency of the patient's condition. Unlike many existing approaches that either simply convert raw signals to symbols using different thresholds, or employ domain-agnostic signal processing methods to summarize sensor data, our technique combines medical knowledge-base with patient-specific personalization to derive symbolic motifs that are both readily communicative of

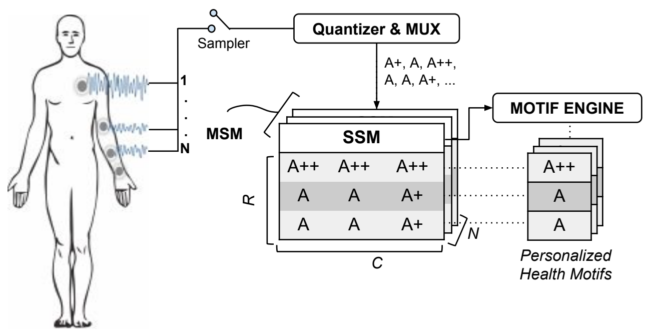

Figure 1: RASPRO framework for transforming the raw sensor data to a patient's health status in the form of Personalized Health Motifs (PHMs).

the patient health status as well as amenable to automated predictive analytics via machine/deep learning techniques. The major contributions as reported in this paper include:

- RASPRO: A personalized medical data summarization technique for converting raw sensor data to "motifs", based on medical domain knowledge.
- Comparative performance analysis (in terms of clinical effectiveness and bandwidth requirement) of RASPRO motifs against a widely accepted symbolic time-series representation technique called Symbolic Aggregate approXimation (SAX).
- Comparative performance analysis of RASPRO motifs against original raw sensor data.

As a concrete application for evaluating the RASPRO framework, we have selected a target disease condition called Acute Hypotensive Episode (AHE). Early detection and prediction of an AHE could greatly reduce the mortality rate.

Experimental results demonstrate the superior clinical accuracy of RASPRO in predicting AHE. In particular, the RASPRO motifs give a maximum F1-score of up to 0.83 as compared to the maximum F1-score of 0.61 given by SAX. Interestingly, RASPRO motifs outperform even the raw sensor data time-series suggesting that "less can be better". We also show that RASPRO motifs provide a much better trade-off between bandwidth cost and efficiency. Due to these advantages, RASPRO motifs could potentially be used to send emergency and critical health status updates to remote doctors even over an SMS without dependence on availability of high bandwidth data networks such as 3G/4G/LTE.

2 RELATED WORK

Data summarization techniques range from simple normal/abnormal threshold-based warning/alert systems such as [2] to more complex multi-sensor data fusion techniques [12] [15]. Machine/deep learning based methods have recently gained attention as potentially powerful frameworks for extracting diagnostic insights from historical data. The literature on this topic is rapidly expanding, but some representative works include [5] [8]. SAX [7] belongs to the class of techniques that convert data into simpler, dimensionally reduced symbolic representations. It is one of the most widely applied techniques for data reduction and has proven to provide very high utility in various domains such as the ones discussed in [3] [9]. The applications of SAX, particularly for healthcare, have been discussed in [16] [6]. We observe that these and many other data reduction/transformation techniques [13] [1] are domain-agnostic. They do not leverage the medical value of the data and thereby potentially miss clinically relevant insights as we will see in the evaluation section (Section 4).

3 PERSONALIZED HEALTH STATUS MOTIFS

We begin this section by describing the RASPRO technique. Let us consider raw time-series data coming in from multiple (N) vital parameter sensors, $S_1, S_2, ..., S_N$, attached to a patient (see Figure 1). There are three major processing steps.

Step 1: The first step is converting the time-series of sampled raw sensor data into a time-series of patient-specific discrete (quantized) severity symbols such as 'A', 'A+', '$A++$', 'A–', '$A--$', etc., where the clinically defined normal range for a given sensor is assigned the symbol "A", while values which are above and below normal ranges are assigned increasing number of "+" and "–" suffixes according to the severity. Unlike typical systems with fixed severity thresholds for a sensor, in RASPRO the number of criticality levels L_{CRIT} and their mapping to corresponding sensor value ranges are customized according to the patient's condition as perceived by the doctor.

Medical knowledge driven decision-making 1: In the case of a patient who has been evaluated in the hospital and is suspected to have mild hypertension, but absolutely no risk for other cardiovascular morbidities, moderate thresholds would be assigned for A+, $A++$, etc., because some fluctuation in Blood Pressure (BP) would not be of concern. On the contrary, a patient with confirmed hypertension whose BP is not brought to target maintenance level, who also has additional cardiac risk factors such as, diabetes, obesity, or hypercholesterolemia, would be assigned more stringent thresholds. Moreover, based on their respective organ risks, a cardiologist, ophthalmologist, or nephrologist would assign different thresholds for the same BP parameter. This essentially makes the criticality range definitions highly personalized.

Step 2: In the next step of the RASPRO system, the time-series of these quantized severity symbols is passed through a multiplexer (MUX) that arranges it into a 2-dimensional single sensor matrix (SSM) of C columns and R rows. Here, each row represents a short continuous burst of observations and between the rows there may be a quiescent period. Multiple SSMs corresponding to different sensors are arranged to form a 3-dimensional Multi-Sensor Matrix (MSM).

Medical knowledge driven decision-making 2: If the doctor wants to monitor the BP of a patient every hour (the BP measurement being taken five times in an hour) for 24 hours, the SSM will have 24 rows with each row containing the five BP values corresponding to that hour. The exact number of columns and rows would be decided by the healthcare professional based upon the patients' condition, sensor type, and diagnostic relevance. The doctor may also decide to increase the sampling rate in case the patient is critical.

Step 3: In the third step of RASPRO, all the quantized severity symbols in a row are temporally summarized to one parameter, called the consensus symbol. A consensus symbol captures the dominant trend in the patient data. All the consensus symbols in an SSM are put together in a column fashion to arrive at a corresponding Personalized Health Motif (PHM) (see Figure 1).

Medical knowledge driven decision-making 3: For the specific purpose of predicting AHE using Mean Arterial blood Pressure (MAP) measurements, in this paper, we define consensus symbol as the most frequently occurring symbol in a row of SSM. If there are multiple candidates for a consensus symbol, the higher severity symbol is selected.

However, doctors might be interested in different kinds of summarization based on the sensor type, diagnostic interest, and patients' health condition. They might, for example, want to know the mean of values, frequency of peaks, value of highest peak, most frequent abnormality or other statistics. Accordingly, RASPRO provides a framework to define summarization differently for various clinical requirements.

In Figure 1, the first row of an SSM is summarized and the corresponding consensus symbol, "$A++$", is determined. A more formal approach to the definition of motifs and other ways of summarizing the data is discussed in one of our earlier works [10] which we do not reproduce here due to space limitation.

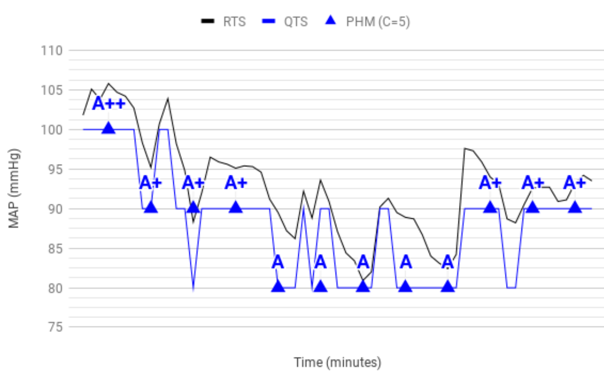

Figure 2: RTS, QTS and MTS of a patient's MAP values over 60 minutes with one MTS symbol every $C = 5$ minutes.

To obtain an intuitive idea of the RASPRO summarization process, we plot in Figure 2, the raw time-series data of a patient's MAP values over a period of 60 minutes (black solid line). We shall refer to this as the raw time-series (RTS). We also do an overlay plot of the corresponding quantized time-series (QTS) represented by the blue solid line. Finally, we also overlay the corresponding RASPRO PHM time-series (which we abbreviate as MTS for Motif time-series), represented by blue triangles together with severity symbols. In the figure, MTS symbols are determined once in every five minutes ($C = 5$).

Medical knowledge driven decision-making 4: Since our target application is episodes of acute hypotension, we consultated cardiologists to make the following criticality levels assignments to the corresponding MAP value ranges: "$A-$": 50-60 mmHg, "A": 60-90 mmHg, "$A+$": 90-100 mmHg, "$A++$": 100-120 mmHg. The also set $C = 5$ and $R = 12$ for a total of 60 minutes of data. These values are used to generate the plots in Figure 2.

In contrast to RASPRO's medical domain knowledge driven decision making process, domain-agnostic summarization techniques do not consider the clinical meaning of MAP values while quantizing. This could, for example, result in a value between 50-70 mmHg being quantized into a single value, hence losing the clinical significance of the data.

4 PERFORMANCE EVALUATION

We conduct three evaluations to analyze the performance of RASPRO: (a) we compare the RASPRO motifs against SAX (one of the most widely used methods for symbolic data summarization) in terms of AHE prediction accuracy, (b) we compare the AHE prediction accuracy of RASPRO against that of raw time-series data, and (c) we compare RASPRO and SAX in terms of their trade-offs between the bandwidth cost of summarization and corresponding predictive power.

4.1 Dataset

For performance evaluation we require vitals sensor data of patients that span a suitably long duration in order to predict the onset of a critical condition, e.g., data from patients who had acute hypotensive episodes (AHE). An AHE is characterized by prolonged decrease in MAP value of a patient below 60 mmHg for more than 90% of the time in a half-hour window. AHE is a potentially fatal condition and needs medication intervention. We obtained AHE patient and control group data from a widely used open source database, MIMIC II [14], which contains multi-parameter records such as MAP, SpO2, ECG, and Non-obtrusive BP. The curated dataset contains at least 60 minutes of MAP data for each subject before the onset of AHE. The dataset contains two groups: 27 patients termed as H-group, who had AHE and another with 15 patients, termed as C-group, who did not experience AHE during their stay in an ICU. Additionally, the dataset also includes the time of onset of an AHE event, marked as T_0. For all the predictive analysis tasks in this paper, we used MAP data that was 60 minutes prior to T_0.

4.2 Symbolic Aggregate Approximation (SAX)

SAX is a widely used technique for data reduction that uses a symbolic representation. In SAX, an arbitrary length of time-series data, n, is reduced to a string with a length smaller than n. This involves mainly two steps: (a) dimensionality reduction using Piecewise Aggregate Approximation (PAA) and (b) discretization using equiprobable symbols.

PAA: In this step, a given time-series data is first amplitude-normalized (by subtracting the global mean and dividing by the global standard deviation) and then partitioned into smaller subsets of size W time steps called frames. A mean value is determined for the data that lies within a specific frame. The frame is then replaced with the computed mean value as its equivalent low-dimensional representation. The above procedure, when repeated for each frame, generates a vector of corresponding mean values, which provides the data-reduced representation of the entire time-series.

SAX: The data-reduced representation obtained after implementing PAA is further discretized to produce symbols of equal probability. The SAX technique assumes that normalized time-series data follow a Gaussian distribution. To allot the symbols, first breakpoints are assigned to produce areas of equal size under the Gaussian curve. Each of these breakpoints is assigned a symbol. Thereafter, the PAA coefficients are mapped to the corresponding breakpoints to generate a symbolic SAX representation which is termed as a

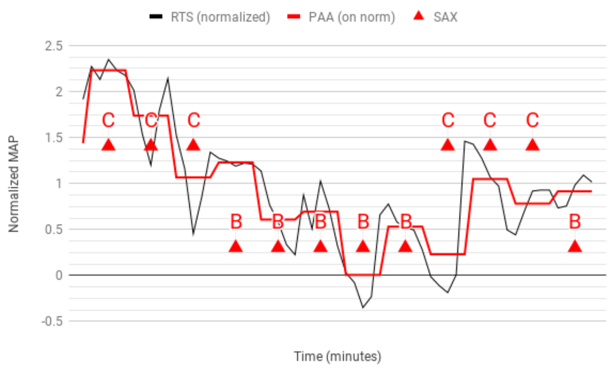

Figure 3: Amplitude-normalized MAP values of a patient over 60 minutes using RTS as well as the corresponding PAA and SAX representation for a frame size of $W = 5$ minutes.

word. We plot in Figure 3, the MAP values in the RTS as well as the corresponding PAA and SAX representations ($W = 5$ minutes). In contrast to MTS (see Figure 2) which incorporates medical domain knowledge, the SAX representation is domain-agnostic (see Figure 3).

4.3 Predictive Power: Comparison with SAX

We compare the performance of RASPRO's MTS against the time-series of SAX symbols on the task of predicting AHEs. For this purpose, 60 minutes of MAP data prior to the onset time T_0 was extracted and transformed to MTS and SAX series. Guided by medical knowledge, we used the following parameter choices to evaluate the performance of MTS: (a) $L_{CRIT} = 7, 9, 13$ levels and (b) $C = 5, 10, 15$ minutes, and the following parameter choices for SAX: (a) alphabet size = 7, 9, and 13 corresponding to values of L_{CRIT} in the RASPRO method and (b) frames of length $W = 5, 10, 15$ minutes corresponding to the values of C for MTS.

For ease of comparison, we use S to denote the alphabet size in SAX and the corresponding value of L_{CRIT} in RASPRO. We use W to denote both the frame length (in SAX) as well as the corresponding value of C (in RASPRO).

The MTS and SAX series were then vectorized[1] and fed into a Support Vector Machine (SVM) binary classifier, which is trained and tested using, respectively, 70% and 30% of the data. We used five fold cross validation and also tried two different kernels (rbf and linear). For each (S,W) pair, we report the best performance from among the two kernels. We chose SVM since it is a widely used classical classification technique which has moderate training and testing complexity and has proved to be highly effective in a number of applications. Other classifiers, e.g., based on neural networks, could potentially provide additional performance improvements.

We measure the predictive power of RASPRO and SAX via the F1-score, which is a statistical measure of binary classification that is the harmonic average of the precision and recall values. Table 1 summarizes the results of the SVM-based prediction task. We observe that RASPRO gives consistently higher F1-scores than SAX across the range of S and W values. The relative F1-score improvements of RASPRO over SAX range from 20.9% for $S = 7, W = 5$ to 93.2% for $S = 13, W = 10$.

	F1-scores					
	$W = 5$ mins		$W = 10$ mins		$W = 15$ mins	
$S =$	RASPRO	SAX	RASPRO	SAX	RASPRO	SAX
7	0.52	0.43	0.61	0.43	0.76	0.47
9	0.77	0.45	0.77	0.47	0.69	0.47
13	0.83	0.61	0.83	0.43	0.83	0.47

Table 1: Comparison of F1-scores for the prediction of an AHE using SAX and RASPRO for different number of symbols S, and summarization time windows, W.

4.4 Comparison with RTS

Additionally, we also compare the F1-score of RASPRO motifs against that of RTS for the same AHE prediction task using SVM. RTS yielded a maximum F1-score of 0.70 compared to the maximum score of 0.83 ($S = 13$ and $W = 15$) for RASPRO motifs which uses only one-fifteenth of the RTS data and that too in quantized form. In fact, except for three (S, W) pairs in Table 1, RASPRO has an F1-score that is higher than 0.70. This clearly demonstrates that RASPRO summarization could be as effective, if not better than the raw time-series representation for many choices of S and W values.

4.5 Bandwidth cost

The trade-off between accuracy of data and bandwidth cost is an important aspect of study when it comes to summarization techniques. We define the bandwidth cost function f as:

$$f = \frac{\log_2 S}{W} \tag{1}$$

where, $\log_2 S$ is the number of bits needed to represent S symbols, and W the size of the summarization window. This cost function captures the data requirement for summarization given a pair of (S, W) values. Figure 4 summarizes our findings. Observe that the F1-score of SAX remains flat (between 0.43–0.47, close to the performance of random guessing) over most bandwidths and increases to 0.61 at $f = 0.74$. In contrast, the F1-score of RASPRO rises roughly monotonically from 0.61 to 0.83 as f increases. The relative improvement over SAX ranges from 20% for $f = 0.56$ bits/unit-time to a maximum of 90% for $f = 0.37$ bits/unit-time. This demonstrates that the medically-informed RASPRO motifs provide much greater clinical value per bit per minute than SAX.

5 DISCUSSION

The RASPRO technique has been designed with the dual objectives of summarization and enhanced interpretability to aid in better and faster clinical decision making. The PHMs present a succinct symbolic representation of a patient's health condition. We also see that these PHMs could be used as input to a machine learning algorithm for a classification task. It is interesting to note that domain-agnostic techniques such as SAX could lead to peculiar interpretations of data. For instance, in the PAA step of SAX it is

[1] Each symbol is assigned a numerical value corresponding to the *lower* end point of the numerical quantization range associated with it.

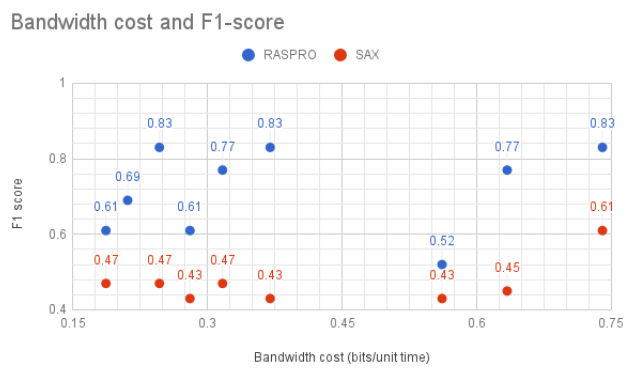

Figure 4: Comparison of bandwidth requirements for RASPRO and SAX using the cost function as defined in Equation (1) shows a clear advantage of RASPRO technique which provides a much higher predictive power for a given bandwidth cost incurred.

possible that multiple abnormal MAP values (say 59 mmHg) and few normal MAP values (say 65 mmHg) could be averaged and summarized to a normal value (above 60 mmgH), hence losing clinical meaning and value. From the experimental results presented in the paper, we observe that by quantizing the sensor data in a principled and clinically-aware way, the predictive power of the data is, in fact, enhanced. The RASPRO motifs showed a maximum F1-score of 0.83 compared to 0.61 for SAX. The performance of RASPRO motifs was also higher compared to the RTS, which gave an F1-score of 0.70. The comparison of bandwidth cost too shows that RASPRO motifs clearly outperform SAX by giving much higher predictive power for every bit used. The better-than-expected results could be attributed to using the medical domain knowledge in the two steps of RASPRO:

- The patient and disease specific quantizer removes irrelevant/redundant data, and represents it in a more clinically meaningful way.
- The motif summarization, when used judiciously with the optimum summarization window, effectively captures the normal and abnormal trends in the patient.

6 CONCLUSION AND FUTURE WORK

In this paper we have presented a novel summarization technique, called RASPRO, which can convert high-volume multi-sensor multi-parameter data from wearable sensors into a patient-specific severity description in the form of symbolic motifs (PHMs). While most of the existing algorithms use domain-agnostic summarization techniques, RASPRO adopts a medical-knowledge-driven approach. In comparison to one of the widely-used techniques for symbolic sequence generation called SAX, we have shown that RASPRO outperforms in improving the clinical effectiveness as represented by the predictive power. We also observed that the bandwidth requirement for sending large amount of sensor data could be reduced using RASPRO. We attribute the efficiency improvement to the use of patient and disease-specific summarization that RASPRO offers.

In the near future, we plan to translate the findings from this study to clinical experiments using the optimum parameters as discovered in our experiments. Additionally, we intend to experimentally validate the RASPRO framework for other disease conditions.

ACKNOWLEDGMENTS

The authors would like to thank Sri. Mata Amritanandamayi Devi, Chancellor, Amrita Vishwa Vidyapeetham for her guidance and support to conduct research that has direct societal impact. We also thank Dr. P Venkat Rangan for his valuable inputs for this paper.

REFERENCES

[1] Insaf Achour, Kaouther Nouira, and Abdelwahed Trabelsi. 2012. Reducing False Alarms in Intensive Care Units Based on Wavelets Technology. *International Journal of Bio-Science and Bio-Technology* 4, 2 (2012), 111–120.
[2] Mirza Mansoor Baig, Hamid GholamHosseini, Martin J. Connolly, and Ghodsi Kashfi. 2014. Real-time vital signs monitoring and interpretation system for early detection of multiple physical signs in older adults. In *Biomedical and Health Informatics (BHI), 2014 IEEE-EMBS International Conference on.* IEEE, 355–358.
[3] Petko Bakalov, Marios Hadjieleftheriou, and Vassilis J Tsotras. 2005. Time relaxed spatiotemporal trajectory joins. In *Proceedings of the 13th annual ACM international workshop on Geographic information systems.* ACM, 182–191.
[4] Jarod T Giger, Natalie D Pope, H Bruce Vogt, Cassity Gutierrez, Lisa A Newland, Jason Lemke, and Michael J Lawler. 2015. Remote patient monitoring acceptance trends among older adults residing in a frontier state. *Computers in Human Behavior* 44 (2015), 174–182.
[5] Eric M Green, Reinier van Mourik, Charles Wolfus, Stephen B Heitner, Onur Dur, and Marc J Semigran. 2017. Abstract 24031: Machine Learning Detection of Obstructive Hypertrophic Cardiomyopathy Using a Wearable Biosensor. *Circulation* 136, Suppl 1 (2017), A24031–A24031. arXiv:http://circ.ahajournals.org/content http://circ.ahajournals.org/content/136/Suppl_1/A24031
[6] Burcu Kulahcioglu, Serhan Ozdemir, and Bora Kumova. 2008. Application of symbolic Piecewise Aggregate Approximation (PAA) analysis to ECG signals. In *17th IASTED International Conference on Applied Simulation and Modelling.*
[7] Jessica Lin, Eamonn Keogh, Stefano Lonardi, and Bill Chiu. 2003. 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.* ACM, 2–11.
[8] Ivo C. Lopes, Binod Vaidya, and Joel J. P. C. Rodrigues. 2013. Towards an autonomous fall detection and alerting system on a mobile and pervasive environment. *Telecommunication Systems* (2013), 1–12.
[9] A McGovern, D Rosendahl, A Kruger, M Beaton, R Brown, and K Droegemeier. 2007. Understanding the formation of tornadoes through data mining. In *5th conference on artificial intelligence and its applications to environmental sciences at the American meteorological society.*
[10] R. K. Pathinarupothi and E. S. Rangan. 2017. Consensus motifs as adaptive and efficient predictors for acute hypotensive episodes. In *2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC).* 1688–1691. https://doi.org/10.1109/EMBC.2017.8037166
[11] Rimma Pivovarov and Noémie Elhadad. 2015. Automated methods for the summarization of electronic health records. *Journal of the American Medical Informatics Association* 22, 5 (2015), 938–947.
[12] J Rajevenceltha, C Santhosh Kumar, and A Anand Kumar. 2016. Improving the performance of multi-parameter patient monitors using feature mapping and decision fusion. In *Region 10 Conference (TENCON), 2016 IEEE.* IEEE, 1515–1518.
[13] T. Rocha, S. Paredes, P. Carvalho, J. Henriques, and M. Harris. 2010. Wavelet based time series forecast with application to acute hypotensive episodes prediction. In *Engineering in Medicine and Biology (EMBC) society, 2010 annual international conference of the IEEE.* IEEE, 2403–2406.
[14] Mohammed Saeed, Mauricio Villarroel, Andrew T. Reisner, Gari Clifford, Li-Wei Lehman, George Moody, Thomas Heldt, Tin H. Kyaw, Benjamin Moody, and Roger G. Mark. 2011. Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): a public-access intensive care unit database. *Critical Care Medicine* 39, 5 (2011), 952. http://dx.doi.org/10.1097/CCM.0b013e31820a92c6.
[15] S. Sreejith, S. Rahul, and R. C. Jisha. 2016. A Real Time Patient Monitoring System for Heart Disease Prediction Using Random Forest Algorithm. In *Advances in Signal Processing and Intelligent Recognition Systems.* Springer International Publishing, 485–500.
[16] Hossein Tayebi, Shonali Krishnaswamy, Augustinus Borgy Waluyo, Abhijat Sinha, and Mohamed Medhat Gaber. 2011. Ra-sax: resource-aware symbolic aggregate approximation for mobile ECG analysis. In *Mobile Data Management (MDM), 2011 12th IEEE International Conference on*, Vol. 1. IEEE, 289–290.