Making Healthcare Sensor Networks Smart

Gavin Wolf, Joseph Prusa, Chancey Kelley, Pablo Pastran and Ray Colucci
Department of Computer and Electrical
Engineering and Computer Science
Florida Atlantic University
Boca Raton, Florida
Email: gwolf2@fau.edu, jprusa@fau.edu, ckelle22@fau.edu, ppastran@fau.edu, rcolucci@fau.edu

***Abstract*—The decreasing cost and increasing capability of small, low-power sensors has enabled the development of practical pervasive continuous healthcare monitoring systems. Such continuous monitoring systems are effective in obtaining and transmitting vital sign information that can be interpreted by medical professionals. In addition, advances in big data analytics and machine learning along with the proliferation of medical data has given rise to a number of useful applications of big data analytics in healthcare. However, utilizing the data from continuous monitoring systems for machine learning presents a unique set of challenges that have not yet been adequately addressed. In this paper, we propose a smart module that can be integrated into existing continuous monitoring systems. Using incremental machine learning algorithms, this module will serve as a diagnostic, predictive and prescriptive tool to assist medical personnel and patients, thereby improving the quality of healthcare delivery.**

***Keywords*—continuous healthcare monitoring, wireless sensor networks, real-time data analytics, incremental machine learning, Internet of Things, smart systems**

I. Introduction

As the world population is aging, life expectancy is increasing and healthcare costs are rising, there is a need for more effective and more efficient healthcare. The percentage of the world's population age 65 and older was 5.1% in 1950, 8.5% in 2015, and is expected to grow to 16.7% by 2050 [1]. In the United States in 2009, an estimated 85.6% of adults 65 and over had one or more chronic conditions, 56.0% had two or more chronic conditions, and 23.1% had 3 or more chronic conditions [2]. Treating chronic diseases accounts for the bulk of healthcare expenditure in the United States: in 2010, 86% of all healthcare spending was for people with one or more chronic conditions [3]. Meeting the need for more effective and more efficient healthcare depends critically on continued innovation in information and communications technologies. Technologies such as wireless sensor networks for continuous patient monitoring, big data analytics and the Internet of Things provide significant opportunities to improve healthcare.

There has been substantial research in developing wireless sensor networks for continuous healthcare monitoring. Such continuous monitoring systems use both wearable and emplaced small, low-power sensors to monitor patients' vital signs. The established continuous monitoring systems are primarily reactive, alerting care providers in the event of a deterioration of vital signs or a detected fall. Building intelligence into continuous monitoring systems to give them the ability to autonomously detect, predict and respond to changes in conditions is a compelling prospect. In addition to more sophisticated emergency alerts, such a system could also aid in the early detection of diseases and the development of personalized treatment plans. However, adding smartness to real-time sensor networks presents a number of challenges, which we argue can be addressed through incremental machine learning algorithms to train streaming or online models.

In this paper, we propose a smart module that can be integrated into existing continuous monitoring systems and we show that the module's proposed algorithms are suitable for a real-time environment. In order for the module to be able to integrate into existing continuous monitoring systems, the module design must be such that the data generated from existing systems can serve as the input for the module and that the module's output can be used to improve the systems' operations. The focus of this paper is on identifying the requirements for a smart module and then proposing a module that meets these requirements. Thus, experiments for validating the module are left for future work.

The rest of this paper is organized as follows. Section II reports related work representing the most prominent continuous monitoring systems. Section III reports demonstrated applications of big data analytics to healthcare. Section IV discusses the types of data collection and aggregation methods used by current continuous monitoring systems. Section V proposes a machine learning module for continuous monitoring systems. Section VI provides our conclusions and future work.

II. Existing Continuous Monitoring Technologies

Advances in wireless sensor network technology have enabled the development of systems that can provide pervasive continuous monitoring. The most critical component of these systems is the wireless body area sensor network, consisting of wearable sensor nodes that can sense, process and communicate a user's vital signs. Since these sensors are networked together, they can communicate the sensed information to a gateway or base station, which can then communicate the information over the Internet to healthcare providers and emergency services. Sensors exist to measure a whole range of vital information including ECG sensors for monitoring heart activity, EMG sensors for muscle activity, EEG sensors for brain electrical activity, blood pressure sensors, motion sensors, and even breathing sensors that can be used to measure a patient's activity level [4]. Two of the most cited research projects related to continuous monitoring are CodeBlue [5] and AlarmNet [6].

The CodeBlue infrastructure consists of wearable vital sign sensors, location beacons, and a suite of protocols and services needed to make the information gathered useful to care providers. CodeBlue uses ad hoc routing techniques to extend the communication range of the devices. Data collection is done through a mobile device like a computer or smartphone. For location tracking, the MoteTrack localization system is used [7]. MoteTrack uses beacon nodes to periodically broadcast messages containing information about the node. CodeBlue is a robust system that encompasses both hardware and software that is scalable. It also uses an ad hoc network design that enables the use of location tracking and multi-hop communication.

The AlarmNet system was designed as an improvement upon CodeBlue. AlarmNet uses both novel and off-the-shelf wireless sensors, an embedded gateway and a back-end database with analytical programs. AlarmNet's programs are designed to provide alerts in the event of an emergency, the definition of which can be defined on a per-patient basis. AlarmNet is designed to be extensible, to accommodate heterogeneous nodes and to be deployable at a large-scale. Other researchers have sought to improve upon AlarmNet by introducing a multitiered device architecture [8]. One big improvement of AlarmNet over CodeBlue is its inclusion of embedded security for private and sensitive data, although some have been critical of its implementation [9].

While the existing continuous monitoring systems, including CodeBlue and AlarmNet, are capable solutions for achieving pervasive monitoring, it is clear that their effectiveness in optimizing patient care could be greatly improved through the inclusion of modern big data analytics technologies, such as the smart module proposed in this paper. The reactive alerts provided by the existing continuous monitoring systems are no doubt useful for responding to emergency situations. Adding intelligence to the systems through real-time machine learning can enable them to make accurate predictions, diagnoses and decisions. A huge opportunity also exists to integrate the data from continuous monitoring systems along with the other sources of medical data, including: Electronic Health Records (EHRs), treatment guidelines, clinical studies, journal articles, medical imaging, pharmacy, insurance, social media posts, and more [10].

III. Machine Learning in Healthcare

Machine learning algorithms are designed to take large amounts of input data and to generate predictions or decisions as output. Computers are said to learn when their decision-making accuracy improves with additional input data. Machine learning algorithms work by learning relationships among a usually large number of features describing the data and a range of potential outcomes. Thus, algorithms can be trained to predict outcomes and provide automated decision-making. While the successful application of machine learning has been seen in a large number of domains, this section will discuss several examples in healthcare. The many successful implementations of machine learning in healthcare give us reason for optimism that a machine learning module for continuous monitoring systems will be similarly successful.

Zhang et al. [11] introduces an algorithm for real-time prediction and diagnosis using EHRs. Their algorithm is referred to as a Very Fast Decision Tree because it uses a fixed amount of memory and can handle a large amount of data in an efficient manner. Their solution is aimed at providing clinical decision support. While their model is designed to operate in real-time, it does so in a rather limited way. Their model can incorporate updates to a patient's EHR and provide decision support in real-time, but it does not support heterogeneous data streams such as vital sign data from wireless sensor networks.

Research on IBM's Watson project proposed and tested a mining system for prognostic decision support using data streams from patients in Intensive Care Unit (ICU) settings [12]. This system can monitor data streams online in real-time, but it also requires offline computation. The offline computation needs are substantial because the system uses cluster learning which requires a large number of distance calculations, demanding a lot of time and computational resources. Their system was shown to be effective at predicting not just the onset of an adverse event, but also patient trajectories over time; however, due to the computational and time requirements it is not suitable for continuous monitoring.

Naive Bayes machine learning algorithms have been used to make personalized risk predictions for future clinical events [13]. Regression algorithms have been successfully applied to high-dimensional data such as MRI brain images [14]. Support Vector Machine (SVM) classifiers have been applied to large medical datasets and image-based classification. Prominent SVM applications include evaluating the link between socioeconomic status and educational attainment [15], and magnetic resonance image-based classification for risks after heart surgery [16].

IV. Real-time collection and processing of wireless sensor data

In regard to the collection and processing of data from sensor networks, the system can be adapted to a wide range of numeric inputs; however, the machine learning algorithms require that the data is at least in a consistent format. For this reason, our proposed system leverages existing infrastructures to minimize implementation costs and increase usability. As a result, the system proposed depends on the aggregation, preprocessing, assembly, and storage of patient data prior to the data being sent to the smart module. This production of data is required because the various systems available are built by different teams and for different operating environments. Wireless sensor networks break down into three major types: cluster, chain, and mesh networks [5] [17]. Each of these types has been utilized successfully in existing continuous

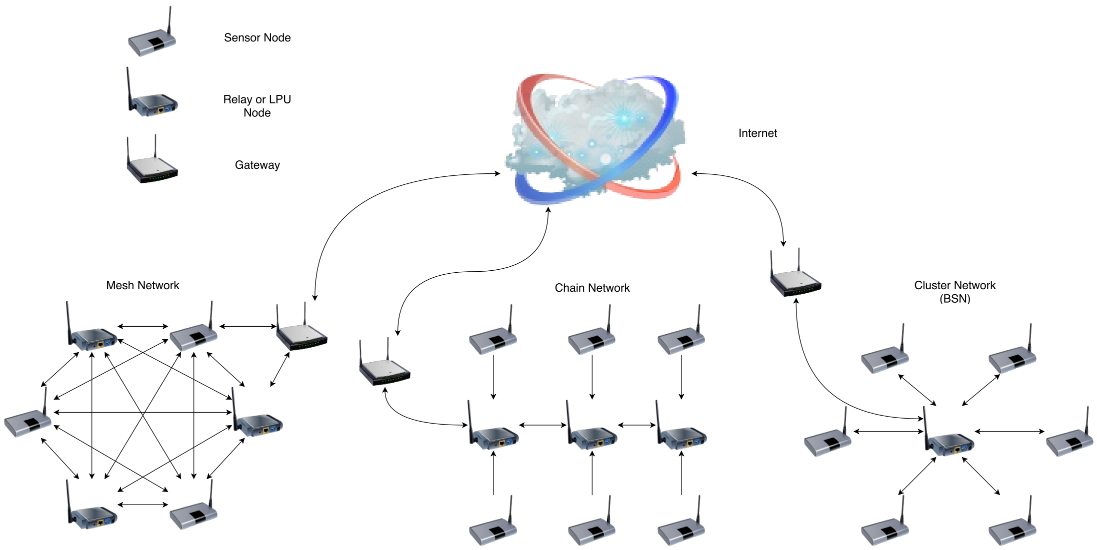

Fig. 1: Network Topology Types

	Type	Noteworthy Features
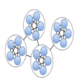	• Cluster	• Self-healing to changes in network topology • Multiple power sources per BSN • Inherent security risk throughout layers
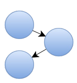	• Chain	• Good sensor node longevity • Requires dedicated nodes to pass messages • Inherent security risk throughout layers
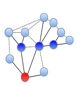	• Mesh	• Very robust transmission paths • Unable to do local processing • Inherent security risk throughout layers

TABLE I: Comparison of Three Network Architectures

monitoring systems. Figure 1 provides an illustration of their architecture.

Cluster networks are characterized by systems such as UbiMon [18]. They have multiple layers of communication by which the data travels, one layer at a time, to reach its final destination. In its native implementation, the data is first collected by a sensor node inside the Body Sensor Network (BSN). The data then travels to the Local Processing Unit (LPU) which processes the data and prepares it for transmission to the gateway. The gateway then sends the data, through the Internet, to the medical servers to be stored and processed. In this manner, the LPU can aggregate multiple sensor nodes needed for the patient and stream the data in batches to the server for later processing and inspection. A benefit of this system is its use of wireless sensors, allowing for flexibility in how to arrange the sensors on a patient's body. Although the BSN being comprised of many individual units has its advantages, it also has the disadvantage of needing multiple power sources per BSN.

Chain networks include systems such as MEDiSN [19]. They are usually characterized by utilizing nodes as the "backbone" of the network. In the case of MEDiSN, the system is broken into two layers of nodes. Sensor nodes collect data from patients and send them to relay nodes. The relay nodes then send that data to the next relay node in the backbone toward the gateway. The sensor node incorporates the medical sensors and the local processing unit onto the same printed circuit board. This sensor node processes, encrypts, and sends the data to relay nodes. As a rule, all sensor nodes do not receive data from other sensor nodes. This allows the sensor nodes to conserve as much power as possible by gating their radios. The relay nodes then transmit the data to the nearest gateway. Since the path is ad hoc in nature, a loss of a relay node will not render the network unusable. Once the data has traveled to the gateway of the network, it travels to the medical server via the Internet. From here it is then processed and made available to medical professionals. This structure has the advantage of sensor node longevity at the expense of node quantity since there are nodes that only pass messages.

An example of a mesh network for continuous monitoring is CodeBlue [5]. Mesh networks are truly ad hoc, and therefore have a decentralized structure where any node has the possibility to talk to any other node by a dynamically generated route. The sensor nodes typically have limited processing power so their function is confined to collecting, preparing and sending data. When the data is ready to be sent, there are three choices: send the data to another sensor node, to a wireless handheld computer, or to a wired, network enabled computer station. Depending on the current network setup and patient information demand, it will automatically establish a route to achieve the desired result. This allows data to travel to a gateway, no matter the scale of the network setup at the time of need. From the gateway, the data travels to the medical server where it is processed and viewed. This implementation of the network structure has the advantage of being very robust in its transmission paths. The implementation has the disadvantage being unable to perform local processing or filtering which could improve battery lifespan.

Between the three types of network structures, there are many similarities in operation. In all network architectures, data is collected, packaged for wireless transmission, and passed to a local gateway. From there is it passed to the medical server for further use. All continuous monitoring systems—regardless of their network structure—have an architecture that includes some sort of medical server to store data. Thus, to achieve compatibility with multiple existing continuous monitoring systems, our proposed smart module is designed to interface with the medical server. The front end of the module would consist of an agent that receives the data packets and deciphers the data's language, type, and bit size. The agent would then do any necessary translation, type casting, and bit scaling to match the machine learning algorithm's input structure. When done in this way, adding new network structures would only require adding a sub-module to the agent. This ability allows for a longer longevity of service life and increasing adaptability to the changing information infrastructure.

V. Proposed Smart Module Design

In this section we present the requirements and design of our proposed machine learning module for continuous monitoring. As this is a flexible system designed to be integrated into existing continuous monitoring systems and this paper is exploratory, this section provides guidelines for selecting appropriate machine learning algorithms and discusses their integration in wireless sensor networks.

There are a wide range of machine learning algorithms that can be used to train predictive models or classifiers for the purpose of making predictions. Though diverse in their operation, these algorithms generally function by learning relationships between features describing instances of data and labels or classes attached to the instance. If we use a time interval of sensor data from a patient as an instance, the features describing the instance are the measurements collected in that interval. Machine learning algorithms can provide automated decision-making to both preemptively anticipate problems and detect problems as they occur based on this data by learning patterns from previous data. Thus, a classifier can be trained from medical data for predictive and diagnostic purposes as an aid to medical personnel. Additionally, the use of machine learning in healthcare has already been shown to be effective, as discussed in Section III.

While machine learning is used in the medical and healthcare domain, current applications are not suitable for continuous monitoring with wireless sensor networks. This specific application has several requirements not encountered when implementing machine learning in other settings. First, classifiers must run in real-time. Second, they must be updatable without having to fully retrain the model. Finally, they must have very high performance in order to maintain patient and doctor trust while minimizing human overhead in managing responses.

Streaming or online classifiers provide the solution to the requirements of being real-time and updatable. These classifiers are variants of machine learning algorithms designed to be trained either incrementally with single instances, or with batches of instances instead of requiring the entire training dataset. The streaming aspect of the model refers to the fact that new instances are used to update the model. As a component of this process, old instances can be forgotten. The rates at which the model learns from new instances and forgets old instances can be controlled by model parameters. Most importantly, during this process the model remains active or online, and can continue to be used for regular operations.

We advocate using incrementally trained classifiers as this is a better match for our real-time goal since batch updates have less frequent update intervals making them less responsive to recent history. While there are numerous incremental classifiers, we suggest the use of either Support Vector Machines or Naive Bayes. Both are popular learners and have been demonstrated to perform well in a wide range of domains. More importantly, both have been successful in the medical domain and both have proven incremental variants.

A. Support Vector Machines

Support Vector Machines (SVMs) [20] construct a hyperplane that divides the instances into two groups while maximizing the distance of the center of each group from the boundary. While the original implementation of SVM separates instances into linearly separable spaces, a non-linear kernel may be adopted. Such a transformation, known as the kernel trick, allows instances to be separated into non-linear spaces and can greatly improve classification performance [21]. An online incremental implementation of SVM has been developed by Tax and Laskov [22]. This implementation of SVM allows for continuous updates, online unsupervised learning and incremental learning, thus allowing the SVM algorithm to be used in a continuous monitoring system. SVM has been demonstrated to perform well in a wide range of classification tasks, including diagnoses of medical conditions such as breast cancer [23], and is a good choice for construction of the machine learning module as it can find complex relationships and is robust to high dimensionality (large numbers of features); however, it does require a large amount of memory to train initially.

B. Naive Bayes

Our second recommended classifier for continuous monitoring is Naive Bayes. Naive Bayes is a learner that used Bayesian probabilities to compute the likelihood of class membership of an instance given its descriptive features. This classifier makes the assumption that all features are independent. While is is often not true, it is still a very effective algorithm [24] and very fast to train and update. However, this assumption means that

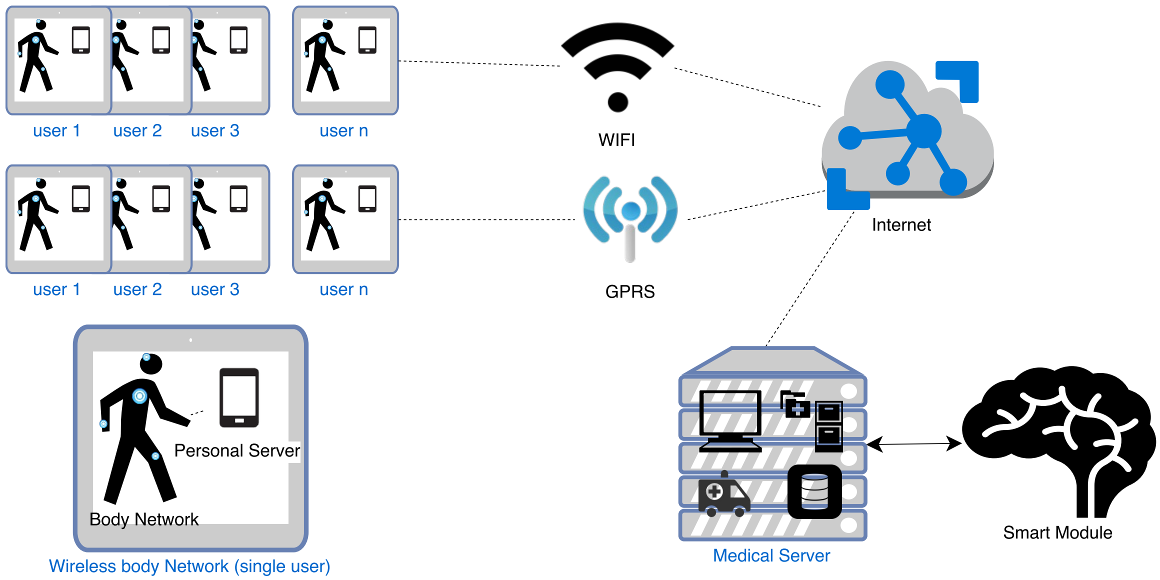

Fig. 2: Illustration of Continuous Monitoring Architecture with Proposed Smart Module

classifiers trained with Naive Bayes do not learn relationships between features. Many implementations exist for training an online incremental Naive Bayes model, we suggest using the online bagging or boosting ensemble approach, described by Oza [25]. These ensemble approaches greatly enhance model performance, while still being extremely fast to train and update compared to batch learners.

C. Machine Learning Module Design

Incremental implementations of SVM and Naive Bayes both meet the algorithm requirements previously outlined. Either machine learning algorithm can be used in the machine learning module and integrated into existing continuous monitoring networks at the medical server level. Within the machine learning module multiple models will be active simultaneously to allow individual models to target specific sets of problems as employing multiple specialized models is more practical than training a single model for everything. Models can be pre-trained and run "out of the box" so that they can be employed immediately. In addition to using sensor data from patients, the module will also leverage EHRs to enhance predictive and prescriptive functionality.

An example of the architecture of a continuous monitoring network with our smart module can be found in Figure 2. When data is sent to the medical server, instances of data organized by a time period and associated with an individual patient will also be sent to the machine learning module for real time analysis and to update the models. The module will feed data into the appropriate models to monitor patient risk for a wide range of diseases and to assess their current state of health. When problems are detected, the patient and medical personnel can be alerted. Additionally, detection of critical and immediate health risks can trigger automatic emergency service deployment. In these situations, emergency personnel will also be provided with detailed information on a patient's history (through access to EHRs) and the effectiveness of various treatment options given the patient's current condition. This better equips emergency personnel and doctors with the information they need to select the correct procedures. Ultimately, this module provides useful diagnostic, predictive and prescriptive decision guidance to patients and medical personnel by using multiple machine learning models to detect patterns in patient data that may not be readily noticeable by human agents.

VI. Conclusions and Future Work

Given the increasing need to provide care for those suffering from medical conditions, along with the rapid pace of technological innovation in wireless sensor networks and big data analytics, there is ample opportunity to use technology to improve the quality of healthcare delivery. Wireless sensor networks in healthcare can provide remote access to a vast amount of real-time patient data, eliminating the barriers of distance and time in tracking patients. Incorporating smart systems using machine learning into these continuous monitoring systems can amplify their range and impact, and can help to make progress in reducing the huge number of preventable diseases and deaths.

In this paper, we propose a smart module that can be incorporated into existing continuous monitoring systems to provide diagnostic, predictive and prescriptive decision support. To lay the foundation for this, we provide a review of existing continuous monitoring sensor networks, current applications of machine learning in healthcare, and how data is aggregated and transported in current systems. Our proposed system uses online, incremental classifiers to provide real-time, continuous and updatable models. Implementations of incremental Naive Bayes or SVM satisfy our design requirements, and either can be chosen depending on specifics of a given module installation. Our machine learning module will operate at the medical server level where it can make use of the incoming stream of patient data to aid medical personnel and patients.

Future work includes developing a system prototype that we can use to run experiments with suitable data sets in order to validate the design and efficacy of the model. We plan to run simulations to test the model. For the dataset, we will simulate a dataset that is informed by the numerous publicly available healthcare datasets. To model the continuous monitoring wireless sensor networks, we will use the ns-3 discrete event network simulator. Once the simulation is up and running, we will develop, test and refine the proposed machine learning algorithms to make decisions and predictions based on the data.

References

[1] C. Bureau, "An aging world," 2015.

[2] CDC, "Percentage of adults age 55 and over with chronic diseases," 2009.

[3] ——, "Chronic disease overview," 2016.

[4] M. Ilyas, S. S. Alwakeel, M. M. Alwakeel, e.-H. M. Aggoune, A. Ambrose, and M. Cardei, "Sensor networks in healthcare," in *Sensor Networks for Sustainable Development*. CRC Press, 2014, pp. 197–214.

[5] D. Malan, T. Fulford-Jones, M. Welsh, and S. Moulton, "Codeblue: An ad hoc sensor network infrastructure for emergency medical care," in *International workshop on wearable and implantable body sensor networks*, vol. 5, 2004.

[6] A. D. Wood, J. A. Stankovic, G. Virone, L. Selavo, Z. He, Q. Cao, T. Doan, Y. Wu, L. Fang, and R. Stoleru, "Context-aware wireless sensor networks for assisted living and residential monitoring," *IEEE network*, vol. 22, no. 4, pp. 26–33, 2008.

[7] K. Lorincz and M. Welsh, "Motetrack: A robust, decentralized approach to rf-based location tracking," in *International Symposium on Location-and Context-Awareness*. Springer, 2005, pp. 63–82.

[8] G. Virone, A. Wood, L. Selavo, Q. Cao, L. Fang, T. Doan, Z. He, and J. Stankovic, "An advanced wireless sensor network for health monitoring," in *Transdisciplinary Conference on Distributed Diagnosis and Home Healthcare (D2H2)*, 2006, pp. 2–4.

[9] P. Gope and T. Hwang, "Bsn-care: A secure iot-based modern healthcare system using body sensor network," *IEEE Sensors Journal*, vol. 16, no. 5, pp. 1368–1376, 2016.

[10] W. Raghupathi and V. Raghupathi, "Big data analytics in healthcare: promise and potential," *Health Information Science and Systems*, vol. 2, no. 1, p. 1, 2014.

[11] Y. Zhang, S. Fong, J. Fiaidhi, and S. Mohammed, "Real-time clinical decision support system with data stream mining," *BioMed Research International*, vol. 2012, 2012.

[12] J. Sun, D. Sow, J. Hu, and S. Ebadollahi, "A system for mining temporal physiological data streams for advanced prognostic decision support," in *2010 IEEE International Conference on Data Mining*. IEEE, 2010, pp. 1061–1066.

[13] J. Wolfson, S. Bandyopadhyay, M. Elidrisi, G. Vazquez-Benitez, D. M. Vock, D. Musgrove, G. Adomavicius, P. E. Johnson, and P. J. O'Connor, "A naive bayes machine learning approach to risk prediction using censored, time-to-event data," *Statistics in medicine*, vol. 34, no. 21, pp. 2941–2957, 2015.

[14] H. Yoshida, A. Kawaguchi, and K. Tsuruya, "Radial basis function-sparse partial least squares for application to brain imaging data," *Computational and mathematical methods in medicine*, vol. 2013, 2013.

[15] S.-M. Zhou, R. A. Lyons, O. Bodger, J. C. Demmler, and M. D. Atkinson, "Svm with entropy regularization and particle swarm optimization for identifying children's health and socioeconomic determinants of education attainments using linked datasets," in *The 2010 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2010, pp. 1–8.

[16] L. P. Kotu, K. Engan, R. Borhani, A. K. Katsaggelos, S. Ørn, L. Woie, and T. Eftestøl, "Cardiac magnetic resonance image-based classification of the risk of arrhythmias in post-myocardial infarction patients," *Artificial intelligence in medicine*, vol. 64, no. 3, pp. 205–215, 2015.

[17] R. Rajagopalan and P. K. Varshney, "Data aggregation techniques in sensor networks: A survey," 2006.

[18] J. W. Ng, B. P. Lo, O. Wells, M. Sloman, N. Peters, A. Darzi, C. Toumazou, and G.-Z. Yang, "Ubiquitous monitoring environment for wearable and implantable sensors (ubimon)," in *International Conference on Ubiquitous Computing (Ubicomp)*. Citeseer, 2004.

[19] J. Ko, J. H. Lim, Y. Chen, R. Musvaloiu-E, A. Terzis, G. M. Masson, T. Gao, W. Destler, L. Selavo, and R. P. Dutton, "Medisn: Medical emergency detection in sensor networks," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 10, no. 1, p. 11, 2010.

[20] C.-W. Hsu and C.-J. Lin, "A comparison of methods for multiclass support vector machines," *IEEE transactions on Neural Networks*, vol. 13, no. 2, pp. 415–425, 2002.

[21] A. Ruiz and P. E. López-de Teruel, "Nonlinear kernel-based statistical pattern analysis," *IEEE Transactions on Neural Networks*, vol. 12, no. 1, pp. 16–32, 2001.

[22] D. M. Tax and P. Laskov, "Online svm learning: from classification to data description and back," in *Neural Networks for Signal Processing, 2003. NNSP'03. 2003 IEEE 13th Workshop on*. IEEE, 2003, pp. 499–508.

[23] M. F. Akay, "Support vector machines combined with feature selection for breast cancer diagnosis," *Expert systems with applications*, vol. 36, no. 2, pp. 3240–3247, 2009.

[24] H. Zhang, "The optimality of naive bayes," *AA*, vol. 1, no. 2, p. 3, 2004.

[25] N. C. Oza, "Online bagging and boosting," in *2005 IEEE international conference on systems, man and cybernetics*, vol. 3. IEEE, 2005, pp. 2340–2345.