

# Taking Meredith out of Grey's Anatomy: Automating Hospital ICU Emergency Signaling

Tarem Ahmed and Supriyo Ahmed  
Department of Electrical and Electronic Engineering  
BRAC University  
Dhaka, Bangladesh  
Email: {tarem, supriyo}@bracu.ac.bd

Fazle Elahi Chowdhury  
Cardiac Intensive Care Unit  
National Heart Foundation Hospital and Research Institute  
Dhaka, Bangladesh  
Email: f-elahi@nhf.org.bd

**Abstract**—In this paper we propose a new algorithm based on kernel machines for automatic, instantaneous detection of emergencies occurring in a hospital Intensive Care Unit. The proposed algorithm takes as input the multitude of vital statistics that are continuously monitored for a critical patient in Intensive Care, learns the underlying pattern between the statistics that is naturally inherent for the particular patient, and instantaneously signals any deviation from this pattern. Through application to real data from a cardiac Intensive Care Unit at a hospital in a developing country, we show that it is possible to easily obtain high detection accuracy with low false alarm rates.

## I. INTRODUCTION

A critical patient in a hospital Intensive Care Unit (ICU) has a large number of probes constantly measuring his or her vital statistics. Round the clock monitoring of these measurements are typically performed by a nurse, with a doctor paged in the event of an emergency. The task of simultaneously monitoring multiple readings on multiple patients becomes tedious and monotonous for a human aide working the long shifts that are typical of the medical profession, specially in understaffed hospitals that are common in developing countries. This consequently increases the risk of important events going unnoticed. Indeed, studies have shown that the optimal concentration span for a human being ranges between 25 and 30 minutes [1].

The well-equipped ICUs in developed countries do have some degree of simple, automated signalling systems installed, with alarms raised when some reading crosses a pre-set critical threshold. These systems are mostly built using simplistic, change point detection algorithms [2]. An example is BioSign<sup>TM</sup> [3], which alerts when major variables deviate  $\pm 3$  standard deviations from their mean values. However, waiting for a critical statistic such as blood pressure or heart rate to exhibit extreme values may be too late for an ICU patient. The set of vital statistics for a patient may together also contain an inherent correlation structure, which may be broken far before any particular statistic displays extreme values. Timely detection of the break of this underlying correlation structure may provide an earlier indication of an impending emergency.

Our goal is the automatic, instantaneous detection of any sudden break in the complicated, underlying structure that is naturally formed by the multitude of biological signals provided by the human body. This inherent correlation structure between the signals may also be unique to a given patient.

As a step towards this goal, we propose the Kernel-based Online Anomaly Detection (KOAD) algorithm [4]. KOAD is an adaptive algorithm that learns the normal pattern in a multivariate timeseries of data, and signals a break in the pattern in real-time. The proposed algorithm is lightweight in terms of computational and memory resources required, and amenable for use with any kind of medical equipment. It is thus ideally suited for application to a hospital in a developing country with limited financial and technological resources. We apply KOAD to timeseries of measurements taken from real cardiac patients at the Intensive Care Unit of the National Heart Foundation Hospital and Research Institute (NHFH&RI), Dhaka, Bangladesh. We show that it is possible to instantaneously detect many subtle events that would otherwise have gone unnoticed and unreported.

### A. Related Work

To the best of our knowledge, this is the first application of a real-time, pattern matching algorithm to automated monitoring of a multidimensional timeseries of observations collected at a hospital ICU [5]. Zhu presented a scheme using Hidden Markov Models (HMMs) in [6] to detect anomalies in blood glucose levels. Her method used historical observations as a benchmark, and required up to 30 days of normal, training data, thereby rendering it ineffective in the time-critical environment of an ICU. Lee et al. have recently proposed an *offline*, block-based approach using Support Vector Machines (SVMs) to classify patients at risk of complications during Percutaneous Coronary Intervention (PCI) [7]. Keogh et al. proposed a data mining technique in [8] to identify *discords* in electrocardiogram (ECG) timeseries data, where a discord is defined as a subsequence with large distance from the rest of the sequence. The technique of Keogh et al. was, however, only a heuristic derived using conventional distance metrics, and the authors' objective was only to obtain a speedup over an earlier brute-force method. The heuristic of Keogh et al. was also restricted to static, single-dimensional data. The follow-up work by Chuah and Fu [9] used an adaptive window-based re-sampling method on the single-dimensional ECG timeseries to improve on the detection accuracies obtained in [8].

### B. Organization of Paper

The rest of this paper is organized as follows. Section II presents our proposed automated emergency detection method using the Kernel-based Online Anomaly Detection (KOAD)

algorithm. Section III presents the results of our experiments conducted on real data collected from patients in a cardiac Intensive Care Unit (ICU). Section IV concludes and outlines the future potential of our approach.

## II. AUTOMATIC EVENT DETECTION

### A. Kernel Functions

Algorithms based on the so-called “kernel trick” involve using a *kernel* function that maps the input data onto a *feature space* of much higher dimension [10], with the expectation that points depicting similar behaviour would cluster in the richer, higher-dimensional space. A suitable kernel function, when applied to a pair of input vectors, may be interpreted as an inner product in the feature space. This subsequently allows inner products in the feature space (inner products of the *feature vectors*) to be computed without explicit knowledge of the feature vectors themselves, by simply evaluating the kernel function:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle \quad (1)$$

where  $\mathbf{x}_i, \mathbf{x}_j$  denote the input vectors and  $\phi$  represents the mapping onto the feature space. Using kernel functions thus allows simple comparison of higher order statistics between the input vectors.

### B. Kernel-based Online Anomaly Detection Algorithm

If the points  $\{\mathbf{x}_t\}_{t=1}^T$  show normal behaviour in the input space, then the corresponding feature vectors  $\{\phi(\mathbf{x}_t)\}_{t=1}^T$  are expected to (also) cluster. Then, it should be possible to explain the region of normality in the feature space using a relatively small *dictionary* of *approximately* linearly independent elements  $\{\phi(\tilde{\mathbf{x}}_j)\}_{j=1}^m$ . Feature vector  $\phi(\mathbf{x}_t)$  is said to be *approximately* linearly dependent on  $\{\phi(\tilde{\mathbf{x}}_j)\}_{j=1}^m$  with approximation threshold  $\nu$ , if the projection error  $\delta_t$  satisfies:

$$\delta_t = \min_{\mathbf{a}} \left\| \sum_{j=1}^m a_j \phi(\tilde{\mathbf{x}}_j) - \phi(\mathbf{x}_t) \right\|^2 < \nu \quad (2)$$

where  $\mathbf{a} = \{a_j\}_{j=1}^m$  is the optimal coefficient vector [11]. Here  $\{\tilde{\mathbf{x}}_j\}_{j=1}^m$  represent those  $\{\mathbf{x}_t\}_{t=1}^T$  that are entered into the dictionary. The size of the dictionary,  $m$ , is expected to be much less than  $T$ , thereby leading to computational and storage savings.

Observe that (2) involves an  $L^2$  norm [12], which may be simplified exclusively in terms of the inner products of  $\phi(\tilde{\mathbf{x}}_j)$  and  $\phi(\mathbf{x}_t)$ , and thus evaluated using the kernel function without explicit knowledge of the feature vectors themselves:

$$\delta_t = \min_{\mathbf{a}_t} \left\{ \mathbf{a}_t^T \tilde{\mathbf{K}}_{t-1} \mathbf{a}_t - 2\mathbf{a}_t^T \tilde{\mathbf{k}}_{t-1}(\mathbf{x}_t) + k(\mathbf{x}_t, \mathbf{x}_t) \right\} \quad (3)$$

where  $[\tilde{\mathbf{K}}_{t-1}]_{i,j} = k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$  and  $[\tilde{\mathbf{k}}_{t-1}(\mathbf{x}_t)]_j = k(\tilde{\mathbf{x}}_j, \mathbf{x}_t)$  for  $i, j = 1 \dots m_{t-1}$ . The optimum *sparsification* coefficient vector  $\mathbf{a}_t$  that minimizes  $\delta_t$  at time  $t$  is then:

$$\mathbf{a}_t = \tilde{\mathbf{K}}_{t-1}^{-1} \cdot \tilde{\mathbf{k}}_{t-1}(\mathbf{x}_t). \quad (4)$$

The expression for error  $\delta_t$  may then be simplified into:

$$\delta_t = k_{tt} - \tilde{\mathbf{k}}_{t-1}(\mathbf{x}_t)^T \cdot \mathbf{a}_t. \quad (5)$$

The Kernel-based Online Anomaly Detection (KOAD) algorithm [4] operates at each timestep  $t$  on a measurement vector  $\mathbf{x}_t$ . It begins by evaluating the error  $\delta_t$  in projecting the arriving  $\mathbf{x}_t$  onto the current dictionary (in the feature domain). This error measure  $\delta_t$  is then compared with two thresholds  $\nu_1$  and  $\nu_2$ , where  $\nu_1 < \nu_2$ . If  $\delta_t < \nu_1$ , KOAD infers that  $\mathbf{x}_t$  is sufficiently linearly dependent on the dictionary, and represents normal behaviour. If  $\delta_t > \nu_2$ , it concludes that  $\mathbf{x}_t$  is far away from the realm of normality and immediately raise a “Red1” alarm to immediately signal an anomaly.

If  $\nu_1 < \delta_t < \nu_2$ , KOAD infers that  $\mathbf{x}_t$  is sufficiently linearly independent from the dictionary to be considered an unusual event. It may indeed be an anomaly, or it may represent an expansion or migration of the space of normality itself. In this case, KOAD does the following: it raises an “Orange” alarm, keeps track of the contribution of the relevant input vector  $\mathbf{x}_t$  in explaining subsequent arrivals for  $\ell$  timesteps, and then takes a firm decision on it.

At timestep  $t + \ell$ , KOAD re-evaluates the error  $\delta$  in projecting  $\mathbf{x}_t$  onto dictionary  $\mathcal{D}_{t+\ell}$  corresponding to timestep  $t + \ell$ . Note that the dictionary may have changed between timesteps  $t$  and  $t + \ell$ , and the value of  $\delta$  at this re-evaluation may consequently be different from the  $\delta_t$  at timestep  $t$ . If the value of  $\delta$  after the re-evaluation is found to be less than  $\nu_1$ , KOAD lowers the orange alarm and keeps the dictionary unchanged.

If the value of  $\delta$  is found instead to be greater than  $\nu_1$  after the re-evaluation at timestep  $t + \ell$ , KOAD performs a secondary “usefulness” test to resolve the orange alarm. The usefulness of  $\mathbf{x}_t$  is assessed by observing the kernel values of  $\mathbf{x}_t$  with  $\{\mathbf{x}_i\}_{i=t+1}^{t+\ell}$ . If a kernel value is high (greater than a threshold  $d$ ), then  $\phi(\mathbf{x}_t)$  is deemed close enough to  $\phi(\mathbf{x}_i)$ . If a significant *number* of the kernel values are high, then  $\mathbf{x}_t$  cannot be considered anomalous; normal traffic has just migrated into a new portion of the feature space, and  $\mathbf{x}_t$  should be entered into the dictionary. Contrarily if almost all kernel values are low, then  $\mathbf{x}_t$  may be concluded to be a reasonably isolated event, and should be heralded as an anomaly. KOAD evaluates:

$$\left[ \sum_{i=t+1}^{t+\ell} \mathbb{I}(k(\mathbf{x}_t, \mathbf{x}_i) > d) \right] > \epsilon \ell, \quad (6)$$

where  $\mathbb{I}$  is the indicator function and  $\epsilon \in (0, 1)$  is a selected constant. In this manner, by employing this secondary “usefulness test”, KOAD is able to distinguish between an arrival that is an anomaly, from one that is a result of a change in the region of normality. If (6) evaluates true, then KOAD lowers the relevant orange alarm to green (no anomaly) and adds  $\mathbf{x}_t$  to the dictionary. If (6) evaluates false, it elevates the relevant orange alarm to a “Red2” alarm.

KOAD also deletes obsolete elements from the dictionary as the region of normality expands or migrates, thereby maintaining a sparse and current dictionary [4]. In addition, it incorporates exponential forgetting [13], [14] so that the impact of past observations is gradually reduced.

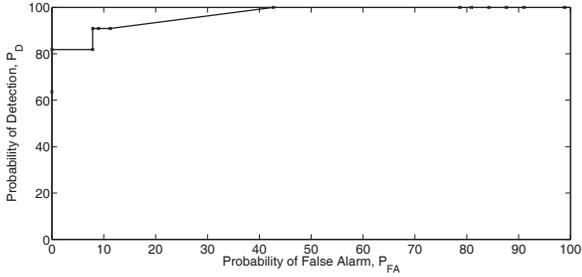


Fig. 1. ROC curve showing KOAD performance. It is clear that close to full detection may be easily obtained with low false alarm rates.

### III. EXPERIMENTS

#### A. Data

Real data was collected from an ICU at the National Heart Foundation Hospital and Research Institute. The data consisted of 100 measurement vectors of 37 vital statistics of 3 patients recorded at 1-minute intervals. Examples of the vital statistics monitored include the heart and pulse rates, the systolic, diastolic and mean blood pressures, the systolic, diastolic and mean pulmonary arterial pressures, the body temperature, the oxygen and carbon dioxide pressures, and important ion concentration levels such as sodium, potassium and chloride. Instances where any of the monitored statistics reached a critical value were manually identified by a cardiac surgeon. It is these labelled instances of particular levels being out of the normal range that we desire instantaneous, automatic alerting to.

#### B. Results

We ran KOAD for a range of values for the thresholds  $\nu_1$  and  $\nu_2$ . The objective was to detect the events that had been manually identified by the surgeon, and obtain as high a detection rate as possible with the least amount of false positives. Figure 1 presents the trade-off between the probability of detection ( $P_D$ ) and the probability of false alarms ( $P_{FA}$ ) as a Receiver Operating Characteristics (ROC) curve. It is obvious from Fig. 1 that it is easy to achieve close-to-full detection with low false alarm rates. This is possible because the normal points *do* indeed cluster in the feature space defined by the chosen kernel mapping, as was initially postulated.

The best values for the thresholds  $\nu_1$  and  $\nu_2$  for a particular application may be determined over a training period using a supervised, cross-validation approach. For our experiments, default values were used for KOAD dropping parameters  $L$  and  $d$ , and the parameters  $\ell$  and  $\epsilon$  that determine orange alarm resolution [4]. As is explained later, that the detection results are not particularly sensitive to the precise setting of any of these supplementary parameters. The Gaussian kernel function [10] proved to be the most suitable.

Figure 2 presents a plot of the KOAD detection statistic  $\delta_t$  for a sample experiment which included 11 known events. Here KOAD was run with thresholds  $\nu_1 = 0.05$  and  $\nu_2 = 0.20$ . The timesteps corresponding to the identified anomalies are shown as red, filled stems. It is clear that most of the identified anomalous timesteps do yield high values for the detection statistic  $\delta_t$  and the instances where critical events occur are easily discernible, a result in agreement with the ROC curve

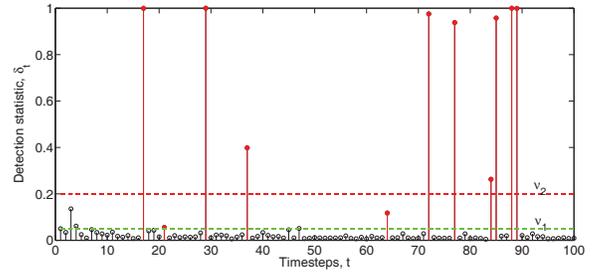


Fig. 2. Progression of the detection statistic  $\delta_t$ , with KOAD run over the full set of 36 monitored vital statistics of a patient across 100 timesteps. Timesteps corresponding to the identified critical events are indicated as red, filled stems. Sample experiment with  $\nu_1 = 0.05$ ,  $\nu_2 = 0.20$ .

depicted in Fig. 1. One of the known anomalies is observed to lie in the “Orange Alarm” range (i.e. between  $\nu_1$  and  $\nu_2$ ) in Fig. 2, while the other identified instances all lie above  $\nu_2$ . The stems in the first few timesteps corresponding to “Green” (normal) cases are seen to be marginally higher than the rest of the normal casw, as the first few timesteps correspond to the training period during which the dictionary is being built.

#### C. Monitoring Specific Vital Statistics

The strength of the KOAD algorithm is that it is able to quickly learn patterns in a streaming sequence of high-dimensional, voluminous data, which is why we advocate its use with the large set of medical measurements typically taken from an ICU patient, to learn the latent patterns and underlying correlation structures which may not be discernible by the doctor’s naked eyes. This said, a doctor may also wish to closely and more precisely monitor a specific measurement, or a specific subset of measurement, for example in the case of an illness which renders this specific subset of medical statistics as the most relevant and crucial. In this subsection we analyse two such cases as examples.

Figure 3 presents a plot of the detection statistic  $\delta_t$ , with KOAD run over a subset of 10 readings (out of the total of 37) corresponding to a patient’s blood pressure levels, for 100 timesteps. Pre-labelling of the data set had indicated 5 instances of specific readings being in dangerous territory, and it is these 5 timesteps that we desire automating alerts to. Oxygen is supplied around the body through blood that is pumped by the heart. Blood enters the heart through the *vena cava* and exits via the *aorta*. Excessive high pressures in the arteries may lead to haemorrhage and cardiac arrests. Timely identification of erratic vascular pressures is crucial to prevent the occurrence of such fatal situations [15].

It is clear from Fig. 3 that all 5 of the identified anomalous timesteps provide high values for the detection statistic  $\delta_t$ , and the instances where critical events occur are easily discernible. Here KOAD was again run with the same thresholds, i.e.  $\nu_1 = 0.05$  and  $\nu_2 = 0.20$ .

Figure 4 presents a plot of the detection statistic  $\delta_t$ , with KOAD run over a subset of 8 readings (out of the total of 37) corresponding to a patient’s various ion concentration level (such as sodium, potassium, chloride, etc) measurements, for 100 timesteps. Pre-labelling of the data set had indicated 7 instances of specific readings being in dangerous territory, and it is these 7 timesteps that we again desire automated

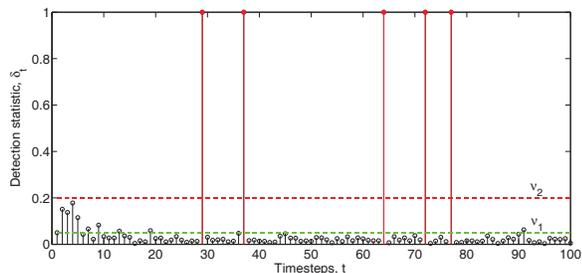


Fig. 3. Progression of the detection statistic  $\delta_t$ , with KOAD run over a set of 10 blood pressures measurements of a patient across 100 timesteps. Timesteps corresponding to the identified critical events are indicated as red, filled stems. Sample experiment with  $\nu_1 = 0.05, \nu_2 = 0.20$ .

alerting to. Ion and electrolyte levels play a vital role in the human body, and a deviation of these levels from the normal range may adversely affect the renal, nervous and respiratory systems. When these ion concentrations deviate out of the normal range, the pH level in the body also changes, which can affect cardiac cell functioning. In such a situation, glycoside drugs need to be immediately injected into the patient to bring the heart rate back to normal and subsequently prevent a heart attack [15].

It is clear from Fig. 4 that all 7 of the identified anomalous timesteps provide high values for the detection statistic  $\delta_t$ , and the instances where critical events occur are easily discernible. Here KOAD was again run with the same thresholds, i.e.  $\nu_1 = 0.05$  and  $\nu_2 = 0.20$ .

#### D. Complexity Analysis

Storage and complexity issues are paramount to real-time applications. In terms of storage requirements, the maximum dimensions of the variables that KOAD stores are  $m \times m$ , where  $m$  is the dictionary size. The computational complexity is  $O(m^2)$  for every standard timestep, and  $O(m^3)$  on the rare occasions when an element removal occurs. KOAD complexity is thus independent of time, making it naturally suited for online use. Our experiments have shown that high sparsity levels are achieved in practice, and the dictionary size does not grow indefinitely. In terms of actual run-time, processing each 37-dimensional timestep took less than one second when run on a laptop computer with Intel i5<sup>TM</sup> processor and standard configuration. This means that it is possible to process multiple readings taken at one-second intervals, a feature which is likely to be very convenient in an ICU environment.

#### E. Parameter Selection

The detection performance of KOAD is primarily a function of the thresholds  $\nu_1$  and  $\nu_2$ . Threshold  $\nu_1$  has the most direct effect on the detection performance, while threshold  $\nu_2$  determines the instant flagging of an anomaly. Our experiments have shown that the performance of a setting remains approximately the same across widely-separated time periods, and optimum settings may be determined after running the algorithm over a training set of labeled data using a supervised, cross-validation approach.

Our experiments have also indicated that the performance is not particularly sensitive to the choice of the orange alarm resolution parameters  $\ell$  and  $\epsilon$  or the dropping parameters  $L$  and

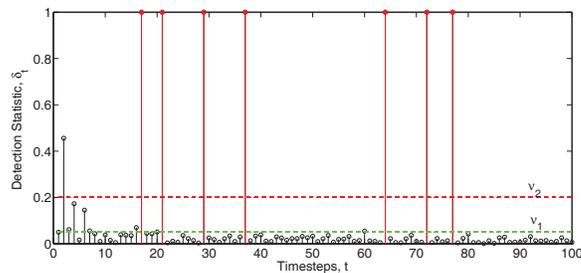


Fig. 4. Progression of the detection statistic  $\delta_t$ , with KOAD run over a set of 8 ion concentration levels of a patient across 100 timesteps. Timesteps corresponding to the identified critical events are indicated as red, filled stems. Sample experiment with  $\nu_1 = 0.05, \nu_2 = 0.20$ .

*d.* They may thus be suited to taste depending on how much of a time-lag is allowable for the orange alarm resolution, for example level of criticality of the particular patient, and the storage resources available to the system.

## IV. CONCLUSION AND FUTURE WORK

In this paper, we have proposed a new algorithm for performing automatic, instantaneous detection of emergencies occurring in a hospital Intensive Care Unit (ICU). The proposed algorithm processes the multitude of vital statistics continuously being monitored for a critical patient in Intensive Care, and learns the underlying pattern between the statistics that is naturally inherent for the particular patient. The algorithm then instantaneously signals a break in this pattern.

The proposed algorithm thus goes beyond the methods presently in place in typical ICUs, which only alert when individual measurements cross pre-set limits. It has run times of the order of hundredths of a second, making it suitable for such a critical environment. It also does not require any expensive or sophisticated components, making it suitable for a hospital in a developing country with financial constraints. Through application to real data from patients in a cardiac ICU at a hospital in a developing country, we have shown that it is possible to easily obtain high detection accuracy with low false alarm rates.

Our future work will expand the testing on data from a larger number of patients, investigate the use of supplementary algorithms to automatically set the KOAD thresholds [16], [17], and the effectiveness of other machine learning algorithms such as the One-Class Neighbor Machine (OCNM) [18] and versions of Principal Component Analysis (PCA) [19], [20] and Kernel Principal Component Analysis (KPCA) [21]. It may be mentioned here that Ahmed *et al.* have done a lot of recent work on adaptive, intelligent pattern matching and anomaly detection [22]–[40], which may also be investigated for this application.

## ACKNOWLEDGMENTS

This project was conceptualized when Tarem Ahmed was affiliated with the research group of Prof. Bernhard Plattner at the Computer Engineering and Networks Laboratory, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland. Data collection and pre-processing was performed by BRAC University students Ferdousi Ahmed, Rashida Bengali, Nushat Mahmud, Amanna Nowshin, Nabila Yeazdani, Fahmida Zahan and Naima Zaman.

## REFERENCES

- [1] C. Donald, "Vigilance," in *People in Control: Human Factors in Control Room Design*, J. Noyes and M. Bransby, Eds. Stevenage, UK: Institution of Engineering and Technology, 2002, pp. 35–50.
- [2] T. Ahmed and R. Rahman, "Survey of anomaly detection algorithms: Towards self-learning networks," in *Security of Self-Organizing Networks: MANET, WSN, WMN, VANET*, A.-S. Pathan, Ed. Auerbach Publications, CRC Press, USA, Sep. 2010, pp. 65–89.
- [3] Biosign Technologies Inc. Company homepage. Accessed: 01 Oct. 2013. [Online]. Available: <http://www.biosign.com/>
- [4] T. Ahmed, M. Coates, and A. Lakhina, "Multivariate online anomaly detection using kernel recursive least squares," in *Proc. IEEE Int. Conf. on Computer Communications (INFOCOM)*, Anchorage, AK, USA, May 2007.
- [5] M. Stanfill, "A systematic review of automated medical coding and classification systems," Master's thesis, Oregon Health & Science University, Portland, OR, USA, May 2009.
- [6] Y. Zhu, "Automatic detection of anomalies in blood glucose using a machine learning approach," *Journal of Communications and Networks*, vol. 13, no. 2, pp. 125–131, Apr. 2011.
- [7] G. Lee, H. Gurm, and Z. Syed, "Predicting complications of percutaneous coronary intervention using a novel support vector method," *J. Am. Med. Inform. Assoc.*, vol. 20, no. 4, pp. 778–786, Jul. 2013.
- [8] E. Keogh, J. Lin, A. Fu, and H. van Herle, "Finding unusual medical time-series subsequences: algorithms and applications," *IEEE Trans. Inf. Technol. Biomed.*, vol. 10, no. 3, pp. 429–439, Jul. 2006.
- [9] M. Chuah and F. Fu, "ECG anomaly detection via time series analysis," *Lecture Notes in Computer Science (LNCS)*, vol. 4743, pp. 125–135, 2007.
- [10] B. Schölkopf and A. Smola, *Learning with Kernels*. Cambridge, MA, USA: MIT Press, Dec. 2001.
- [11] Y. Engel, S. Mannor, and R. Meir, "The kernel recursive least squares algorithm," *IEEE Trans. Signal Proc.*, vol. 52, no. 8, pp. 2275–2285, Aug. 2004.
- [12] W. Rudin, *Functional Analysis*, 2nd ed. New York, NY, USA: McGraw-Hill, Jan. 1991.
- [13] A. Sayed, *Fundamentals of Adaptive Filtering*. Hoboken, NJ, USA: Wiley, Jun. 2003.
- [14] S. Haykin, *Adaptive Filter Theory*, 4th ed. Upper Saddle River, NJ, USA: Prentice-Hall, Sep. 2001.
- [15] N. Colledge, B. Walker, and S. Ralston, *Davidson's Principles and Practice of Medicine*, 21st ed. London, U.K.: Churchill Livingstone, Mar. 2010.
- [16] T. Ahmed, "Online anomaly detection using KDE," in *Proc. IEEE Global Communications Conf. (GLOBECOM)*, Honolulu, HI, USA, Nov. 2009.
- [17] T. Ahmed, B. Oreshkin, and M. Coates, "Machine learning approaches to network anomaly detection," in *Proc. ACM/USENIX Workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SysML)*, Cambridge, MA, USA.
- [18] T. Ahmed, X. Wei, S. Ahmed, and A.-S. K. Pathan, "Automated intruder detection from image sequences using minimum volume sets," *International Journal of Communication Networks and Information Security*, vol. 4, no. 1, pp. 11–17, Apr. 2012.
- [19] X. Wei, J. Fan, M. Chen, T. Ahmed, and A.-S. K. Pathan, "SMART: A subspace based malicious peers detection algorithm for P2P systems," *International Journal of Communication Networks and Information Security*, vol. 5, no. 1, pp. 1–9, Apr. 2013.
- [20] X. Wei, T. Ahmed, M. Chen, and A.-S. K. Pathan, "PeerMate: A malicious peer detection algorithm for P2P Systems based on MSPCA," in *Proc. IEEE Int. Conf. on Computing, Networking and Communications (ICNC)*, Lahaina, HI, USA, Jan. 2012.
- [21] T. Ahmed, X. Wei, S. Ahmed, and A.-S. K. Pathan, "Efficient and effective automated surveillance agents using kernel tricks," *SIMULATION: Transactions of The Society for Modeling and Simulation International*, vol. 89, no. 5, pp. 11–17, May 2012.
- [22] T. Ahmed, A.-S. K. Pathan, and S. Ahmed, "Learning algorithms for anomaly detection from images," *IGI Global Int. J. of System Dynamics Applications*, vol. 4, no. 3, pp. 43–69, Jul. 2015.
- [23] —, "Adaptive algorithms for automated intruder detection in surveillance networks," in *Proc. IEEE Int. Conf. on Advances in Computing, Communications and Informatics (ICACCI)*, Delhi, India, Sep. 2014.
- [24] T. Ahmed, S. Ahmed, and A.-S. K. Pathan, "Automated surveillance in distributed, visual networks: An empirical comparison of recent algorithms," *Int. J. of Control and Automation*, vol. 7, no. 3, pp. 389–400, Mar. 2014.
- [25] —, "Automated surveillance in distributed visual networks: A comparative study," in *Int. Workshop on Mobile and Wireless (WMOBILE/Wireless)*, Jeju, Republic of Korea, Dec. 2013, *Advanced Science and Technology Letters*, vol. 42 (Mobile and Wireless 2013), pp. 1–4.
- [26] T. Ahmed and A.-S. K. Pathan, "Automated surveillance approaches," in *International Islamic University Malaysia (IIUM) Postgraduate Students Colloquium*, Kuala Lumpur, Malaysia, Oct. 2013, abstract, poster and presentation.
- [27] T. Ahmed, "Efficient and effective automatic surveillance approaches," Ph.D. dissertation, International Islamic University Malaysia (IIUM), Kuala Lumpur, Malaysia, Sep. 2013.
- [28] T. Ahmed, X. Wei, S. Ahmed, and A.-S. K. Pathan, "Automated visual surveillance using kernel tricks," in *INRIA/ENS Computer Vision and Machine Learning Summer School (CVML)*, Paris, France, Jul. 2013.
- [29] T. Ahmed and A.-S. K. Pathan, "An effective automated surveillance system," in *International Islamic University Malaysia (IIUM) Postgraduate Students Colloquium*, Kuala Lumpur, Malaysia, Dec. 2012, abstract, poster and presentation.
- [30] —, "Towards automated surveillance systems," in *International Islamic University Malaysia (IIUM) Postgraduate Students Colloquium*, Kuala Lumpur, Malaysia, Feb. 2012, abstract, poster and presentation.
- [31] T. Ahmed, X. Wei, S. Ahmed, and A.-S. K. Pathan, "Intruder detection in camera networks using the one-class neighbor machine," in *Proc. American Telecommunications Systems Management Association (ATsMA) Networking and Electronic Commerce Research Conf.*, Riva del Garda, Italy, Oct. 2011.
- [32] T. Ahmed, "Online anomaly detection using adaptive, learning algorithms," in *Traffic Monitoring and Analysis (TMA) PhD School*, Naples, Italy, Jun. 2011, poster and presentation, organizers' travel grant winner.
- [33] T. Ahmed, Sabrina Ahmed, Supriyo Ahmed, and M. Motiwala, "Real-time intruder detection in surveillance systems using adaptive kernel methods," in *Proc. IEEE Int. Conf. on Communications (ICC)*, Cape Town, South Africa, May 2010.
- [34] T. Ahmed and M. Coates, "Machine learning algorithms for anomaly detection in optical networks," in *Workshop on Optimization of Optical Networks (OON)*, Montreal, QC, Canada, May 2007, abstract and poster.
- [35] —, "Machine learning algorithms for anomaly detection in agile all-photonics networks," in *Agile All-Photonic Networks AAPN Annual Research Review*, Ottawa, ON, Canada, Jun. 2007, abstract and poster, runner-up award winner.
- [36] T. Ahmed, B. Oreshkin, and M. Coates, "Traffic incident detection from road camera networks," in *Centre for Advanced Systems and Technologies in Communications (SYTACom) Research Workshop*, Quebec City, QC, Canada, May 2007, abstract and poster, award nominee.
- [37] T. Ahmed and M. Coates, "Online anomaly detection for optical networks," in *AAPN Annual Research Review*, Ottawa, ON, Canada, Jun. 2006, abstract and poster, award nominee.
- [38] —, "Online anomaly detection for optical networks," in *Workshop on Optimization of Optical Networks (OON)*, Montreal, QC, Canada, Apr. 2006, abstract and poster.
- [39] —, "Predicting flow rate densities," in *Agile All-Photonic Networks (AAPN) Annual Research Review*, Ottawa, ON, Canada, Jun. 2005, abstract and poster.
- [40] T. Ahmed, N. Saberi, and M. Coates, "Time-slot reservation in all-photonics networks based on flow prediction," in *Workshop on Optimization of Optical Networks (OON)*, Montreal, QC, Canada, Apr. 2005, abstract and presentation.