Efficient and Low Latency Detection of Intruders in Mobile Active Authentication

Pramuditha Perera, *Student Member, IEEE* and Vishal M Patel, *Senior Member, IEEE*

***Abstract*—Active authentication (AA) refers to the problem of continuously verifying the identity of a mobile device user for the purpose of securing the device. We address the problem of quickly detecting intrusions with lower false detection rates in mobile AA systems with higher resource efficiency. Bayesian and Minimax versions of the Quickest Change Detection (QCD) algorithms are introduced to quickly detect intrusions in mobile AA systems. These algorithms are extended with an update rule to facilitate low frequency sensing which leads to low utilization of resources. Effectiveness of the proposed framework is demonstrated using three publicly available unconstrained face and touch gesture-based AA datasets. It is shown that the proposed QCD-based intrusion detection methods can perform better than many state-of-the-art AA methods in terms of latency and low false detection rates. Furthermore, it is shown that employing the proposed resource efficient extension further improves the performance of the QCD-based setup.**

***Index Terms*—Biometrics, continuous authentication, active authentication, fastest detection.**

I. Introduction

MODERN smartphone devices have revolutionized the consumer lifestyle in the span of the last decade. With the services they provide in communication, networking, entertainment, education, finance and even in navigation, it has become the indispensable human companion. As a result, modern mobile devices contain agglomeration of personal user information ranging from personal photographs, contacts, banking information to passwords. Industry surveys show that 10% of phone theft victims claim to have lost confidential information, 9% of the victims have experienced identity theft, and 12% of the victims have experienced fraudulent charges on their account [1]. The total cost associated with information theft is substantial considering that 2.1 million cases of phone theft were reported in 2015 in the U.S. alone [2]. Therefore, modern mobile devices hold a substantial intangible value in terms of content in addition to their physical value. In this context, mobile user authentication is paramount in safeguarding information security of the user. However, developing such a system has proved to be a challenging task due to the high expectations mobile users have in terms of usability. Mobile users prefer mobile protocols with minimum interference and mobile services with minimal impact on battery life and processing speed of the device [3].

Existing methods for authenticating users on mobile devices are based on passwords, PINs, secret patterns or fingerprints. Recent studies have shown that about 34% of smartphone users in the U.S. do not even secure their mobile devices with passwords or PINs [4], [5], [6], [7]. This is mainly due to the fact that users find entering passwords and PINs on virtual keyboards time consuming, cumbersome and error-prone [8]. Vulnerability of swipe pattern-based passwords has also been widely discussed in the literature. For instance, swipe patterns can be decoded based on the smudge patters on the touchscreen [9] or through shoulder surfing attacks, where an attacker attempts to memorize the swipe pattern through line-of-sight [10]. Furthermore, as long as the mobile device remains active, typical devices incorporate no mechanisms to verify that the user originally authenticated is still in control of the mobile device. Thus, unauthorized individuals may improperly obtain access to personal information of the user if a password is compromised or if a user does not exercise adequate vigilance after initial authentication.

In order to overcome these issues, both biometrics and security research communities have developed techniques for Active Authentication (AA) on mobile devices. These methods essentially make use of the physiological and behavioral biometrics using built-in sensors and accessories such as gyroscope, touchscreen, accelerometer, orientation sensor, and pressure sensor to continuously monitor the user identity. For instance, physiological biometrics such as face can be captured using the front-facing camera of a mobile device and can be used to continuously authenticate a mobile device user [11], [12], [13], [14], [15], [16]. On the other hand, sensors such as gyroscope, touchscreen and accelerometer can be used to measure behavioral biometric traits such as gait, touch gestures and hand movement transparently [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29]. Some of the other CA methods are based on web browsing behavior [30], behavior profiling [31], texting pattern [32], [33], and body prints [34]. In particular, it has been shown that multimodal biometrics-based CA methods perform better than unimodal systems [11], [35], [36]. Note that the terms continuous authentication [35], [37] active authentication [38], implicit authentication [39], [40], and transparent authentication [41] have been used interchangeably in the literature.

It is well known that a balance needs to be made between security and usability of a biometrics-based AA system [5], [42], [43]. The design of usable yet secure AA systems raises crucial questions concerning how to solve conflicts between mobile security and usability. In order to balance usability and security of an AA scheme, we must address the following fundamental challenges.

1. Accuracy : How accurately does a mobile AA system detect an attacker or an intruder? Due to limitations of

Pramuditha Perera is with the Department of Electrical and Computer Engineering, Rutgers, The State University of New Jersey, New Jersey, USA. Email: pramuditha.perera@rutgers.edu

Vishal M. Patel is with the Department of Electrical and Computer Engineering, Rutgers, The State University of New Jersey, New Jersey, USA. Email: vishal.m.patel@rutgers.edu.

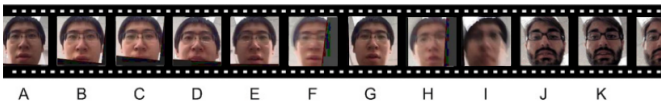

Fig. 1: The problem of quick intrusion detection in face-based AA systems. (A-I) show the genuine user with varying facial expressions. An intrusion occurs starting from (J). Active authentication systems should be able to detect intrusions as quickly as possible without causing too many false detections.

representation and classification models on mobile devices, behavioral and physiological biometrics-based methods do not provide good accuracy in practice [35], [3]. The AA system will be of little use in terms of security if it produces a high degree of false positives. On the other hand, a higher false negative rate would severely degrade the usability of the technology. Many recent approaches in the literature have attempted to address this factor by proposing better features and classifiers [35].

2. Latency : How long does it take to detect an attacker? If an AA system takes too long (e.g. 1-3 minutes) to detect an intrusion, it would grant an intruder plenty of time to extract sensitive information prior to the lock down. Hence, unless intruder detection is sufficiently fast, the AA system would hold a little value in practice no matter how high its detection accuracy is.

Consider a series of observations captured from a front-facing camera of an Android device shown in Figure 1. Frames (A-I) belong to the genuine user of the device. From frame J onwards an attacker starts to operate the device. In this scenario, frame J signifies a change point (i.e. an intrusion). The AA system should be able to detect intrusions with a minimal delay while maintaining a low rate of false detections. For instance, note the changes in genuine user's images in frames (D-F) due to camera orientation and facial expressions. While having a fast response, an AA system ideally should not falsely interpret these variations as intrusions.

3. Efficiency : How much resource does the system use? By definition, mobile AA systems are continuous processes that run as background applications. If they consume considerable amount of resources, memory and processing power, it could slow down other applications and cause the battery to drain quickly. Despite the improvements in mobile memory and processors, battery capacity remains to be a constraint due to limitations in heat transfer and space [44]. Therefore, it can be expected to be the bottleneck in terms of efficiency in years to come. If an AA application causes battery to drain too quickly, then it is unrealistic to expect the users to use AA technology as they would typically opt out from using such applications [45]. Therefore, efficiency has a huge impact over the usability of AA as a technology. Recently, [46] studied the efficiency of a mobile AA system based on face biometric. Experiments were conducted on a Google Nexus 5 device with 2GB of RAM and a quad core 2.2GHz CPU. It was shown that the normal usage of the device consumes about 520 mW of power and the facial attribute-based AA framework running at 4 frames per second consumes about 160.8mW additional power. It is needless to say that nearly 30% increase in power consumption would take a toll on battery duration. A trivial solution for this problem would be to decrease the sampling rate of data acquisition. However, effects of such a measure on the detection performance have not been studied in the literature.

Many existing AA systems attempt to improve the accuracy of the system by proposing sophisticated features and classifiers. However, how fast an AA system could detect an intruder has not been widely studied in the literature. Yet, it remains to be an important feature of an AA system. In this paper, we address the problem of quickly detecting intrusions with lower false detection rates in mobile AA systems. We propose Quickest Change Detection (QCD), which is a well-studied problem in statistical signal processing and information theory, for the purpose of intrusion detection in mobile AA systems. Figure 2 gives an overview of the proposed method. As opposed to a conventional AA system, the proposed system utilizes all past observations along with distributions of match and non-match data of the genuine user to arrive at a decision. The proposed method does not require a specific feature nor a specific classifier; therefor it can be built upon any existing AA system to enhance its performance.

This paper makes the following contributions.

1) We propose Bayesian and Minimax versions of the QCD algorithms for fastest detection of intrusions in mobile AA systems.
2) We propose a sampling scheme with a rule on score update to increase the efficiency of the proposed QCD process.
3) Efficient data driven QCD algorithms are proposed for fastest intruder detection.
4) ADD-PFD curve is proposed as a qualitative performance measure for comparing performances in quickest intrusion detection problems.

A preliminary version of this work appeared in [47], which describes just the Bayesian and minimax QCD algorithms for AA. Data efficient QCD algorithms and extensive experimental evaluations are extensions to [47].

Remainder of the paper is organized as follows. A brief introduction to related work in AA is given in Section II. Section III gives an overview of tradition AA pipeline and describes a new measure for evaluating different intruder detection algorithms. In Section IV, we first briefly introduce the idea behind QCD and then present our proposed framework for QCD-based intrusion detection in mobile AA systems. In Section V, we introduce data efficient QCD methods which can be used to improve the efficiency of the proposed system. Experimental results are presented in Section VI and Section VII concludes the paper with a brief summary and discussion. Notations and abbreviations used in the paper are listed in Tables I and II, respectively.

II. Related Work

In this section, we review some related works on AA. [25] proposed the usage of the accelerometer sensor to capture the behavioral information about the user for continuous

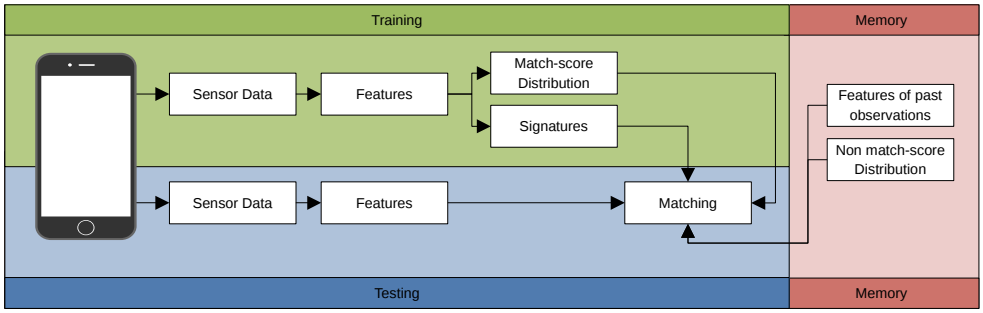

Fig. 2: An overview of the proposed QCD-based AA method.

TABLE I: Notations used in this paper.

Notation	Definition
x_i	Match score obtained at the i^{th} time instance
f_0	Density of matched scores
f_1	Density of non-matched
$E[.]$	Expectation operator
$P[.]$	Probability function
$(x)^+$	Positive portion of x
T	Time at an intrusion occurs
π_n	Probability of intrusion occurring at time n
ρ	Probability of an intrusion occurring
C_α	Set of possible solutions for threshold α
p_n	Probability of change has occurred at time n
$L(.)$	Likelihood ratio
M_i	Indicator of whether observation i is recorded

TABLE II: List of abbreviations.

Abbreviation	Meaning
AA	Active authentication
ADD	Average detection delay
ANO	Average number of observations
APO	Average percentage of observations
BQCD	Bayesian Quickest change detection
CDC	Change duty cycle
CPU	Central processing unit
E-BQCD	Efficient Bayesian quickest change detection
E-MQCD	Efficient minimax quickest change detection
FAR	False acceptance rate
LBP	Local binary pattern
MQCD	Minimax Quickest change detection
PFD	Probability of false detections
PIN	Personal identification number
QCD	Quickest change detection
RAM	Random access memory
WADD	Worst average detection delay

authentication. Based on the same rationale, [26] proposed using the inbuilt sensors of a phone to estimate the gait of the user to perform authentication. It was shown in [27] that even the way how a user holds the phone can be discriminative for the purpose of AA.

In [17], the authors proposed a framework where AA is performed based on the swiping patterns of the users. In their work they introduced a 31 dimensional feature for this purpose. In [21], a new two dimensional feature was introduced for touch gesture-based AA. Sparse dictionary-based classification method is employed with linear and non-linear kernels in [22] to achieve more robust user identification. It was shown in [19] that utilization of external digital sensors can result in even better AA systems. A comparison of various touch-based AA systems on a new dataset is presented in [18].

Most of face-based AA systems are based on hand crafted features due to limited processing power a device holds. In [14], Adaboost classifier and LBP feature are used for face detection and face authentication in mobile devices. In [13], a facial attribute-based continuous face authentication was proposed for AA. A domain adaptive sparse dictionary-based AA system was proposed in [28], by projecting observations of different domains into a common subspace through an iterative procedure. [11] proposed to fuse face and voice for obtaining more robust AA. In [15], face modality was fused with gyroscope, accelerometer, and magnetometer modalities for more robust authentication.

Some of the recent works in AA focus on alternative modalities for performing authentication. Work in [23] focuses on micro movements of the users' touch gestures whereas [24] carries out behavioral profiling based on stylometry, GPS location and web browsing patterns. Please see [35] for more detailed survey of various AA methods.

We note that all these works focus on either developing a new AA modality or improving performance of prevailing modalities by introducing better features or classifiers. None of these works have addressed latency in decision making in the event of an intrusion; nor have they provided a discussion about the efficiency of their methods. Therefore, in terms of the fundamental challenges in AA, these methods do not address challenges of *latency* and *efficiency* which are paramount for a functioning AA system. Intruder detection scheme proposed in [48] uses a single step time series along with a time decaying function to update the belief on the user's authenticity. A false accept rate (FAR) based mapping function is summed with a decaying function in [15] to produce a score which is used to detect intrusions through thresholding. All of these methods are ad-hoc in nature and do not generalize well to more practical and generic setups.

III. Intruder Detection in AA

A typical AA system consists of several stages as illustrated in Figure 3. Initially, sensor data of the genuine user is obtained through an enrollment phase and a set of features

are extracted from the enrolled data. Face images, swipe gesture coordinates, gyroscope/ accelerometer readings and microphone amplitudes are popular choices of data for this purpose. These set of features serve as the gallery at the matching stage.

Upon the initial login of the user, the device continuously collects the same set of data as before during the normal operation of the device. This stage is the Data Acquisition phase shown in Figure 3. Features generated with the collected data are compared against the gallery using a Biometric System using a suitable authentication algorithm. At the end of the comparison phase, a match score x_i is obtained. At the n^{th} time instance based on previously observed matched scores $x_1, x_2, ..., x_n$, a decision is made as to whether an intrusion has occurred or not. If an intrusion has occurred, the phone is locked and the user is prompted to verify his/her identity by the means of a primary verification method. This typically takes the form of a password or a primary biometric such as fingerprint. Otherwise, the user is allowed to continue with the device until the next sensor observation.

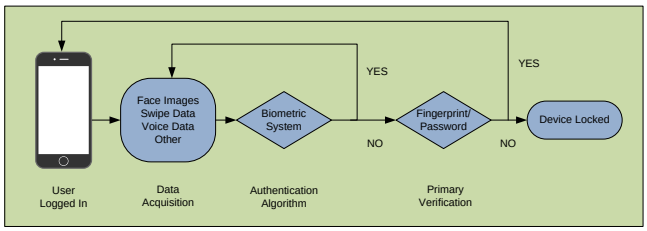

Fig. 3: An overview of a typical AA system.

The score distribution obtained as explained for the genuine user is henceforth referred to as the match score distribution (f_0). Similarly, score distribution of non-genuine users (intruders or attackers in this context) is referred to as the non-match distribution (f_1). Hence, when an intrusion occurs, the distribution of observations changes from being match to non-match. Therefore, an intrusion point is treated as a change point. With this background, we use the words pre-change distribution and match distribution interchangeably. Similarly, post-change distribution and non-match distribution are used interchangeably. If the match distribution has considerable overlap with the non-match distribution, then the detection results tend to be poor. This is typically the case in mobile devices where sensor data acquisition appears in an unconstrained setup. For example, in the case of face-based AA, face images captured by the front-facing camera contain profile faces, tilted faces as well as partial faces. Therefore, the resulting match score distribution tends to be broad. On the other hand, usage of more sophisticated tools that provides better separation between the two distributions are not preferred for mobile applications due to hardware limitations of the device. As a result, match and non-match distributions tend to overlap considerably. In this context, a more scientific approach backed by a theoretical reasoning is essential to perform the detection of the change. To this end, we propose the use of statistical QCD to detect intrusions in the mobile AA systems.

In the following subsections, we identify two essential characteristics such an AA system needs to possess in order for it to be useful in practice.

A. Average Detection Delay (ADD)

The primary goal of an AA system is to promptly detect intrusion when the intruder attempts to access the device. Therefore, detection delay of intruder attempts is an important characteristic of a mobile AA system. If the system requires large number of sensor samples to identify an intrusion, there is a possibility that information theft has already occurred by the time intrusion was detected. Hence, from the point view of security [49], it is more desirable to have an AA system with a low intrusion detection delay.

B. Probability of False Detections (PFD)

On the other hand, if an AA system generates large number of false intruder detections, it would reduce the usability [49] of the user. For example, consider the system shown in Figure 3. The AA system prompts the user to enter a password every time AA fails. If the AA system consistently generates false intruder detection alarms, the user will be prompted to enter the password regularly - thereby greatly degrading consumer experience (usability).

As a consequence, Average Detection Delay (ADD) and Probability of False Detections (PFD) play a vital role in any AA system. If T is the real change point, mathematically ADD and PFD at time τ are defined as follows

$$\begin{aligned} ADD(\tau) &= E[(\tau - T)^+] \\ PFD(\tau) &= P[\tau < T], \end{aligned} \tag{1}$$

where $E[.]$ and $P[.]$ are the expectation and probability with respect to τ, respectively and $[(x)^+]$ denotes the positive part of x.

From these definitions, one can see that there is an inverse correlation between these two quantities. Generally, obtaining more sensor samples enhances the chance of making a more accurate decision on whether an intrusion has occurred or not. However, this can only be done at the cost of having a relatively larger intrusion detection delay. Therefore, there is always a trade-off between intrusion detection delay and false intruder detection rate. Since, the relationship between ADD and PFD characterizes the performance of an AA system, we propose using the ADD-PFD graph as a tool to compare the performance of different AA systems. Shown in Figure 4 are a set of ADD-PFD plots drawn for practical non-sequential AA systems. As expected, in order to obtain very accurate detections (corresponding to a lower PFD), more samples are required to be processed. Moreover, according to Figure 4, making a decision based on fewer samples are prone to more false intruder detections.

Adhering to security and usability principles [49], the objective of an AA system is to be able to detect intrusions while ensuring probability of false intruder detection is very low. Therefore, the AA system should operate in a region where both ADD and PFD are comparatively low. For example, for

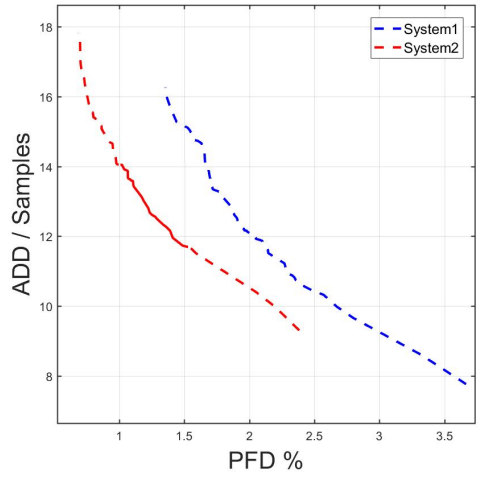

Fig. 4: PFD-ADD curves of two AA systems. Out of the two, the system corresponding to the red curve has better performance. One of the possible operable regions for a practical system is shown in a solid red line.

the system represented by the red curve in Figure 4, a practical choice would be to operate in the region denoted by the solid line.

Based on this rationale, a better AA system should have a PFD-ADD curve operating below other comparable systems. For example, considering the two operating curves shown in Figure 4, system corresponding to the red colored line has a better performance since its operating curve lies at a lower space compared to the other system.

IV. Quickest Change Detection

Quickest Change Detection is a branch of statistical signal processing that thrives to detect the change point of statistical properties of a random process [50], [51], [52]. The objective of QCD is to produce algorithms that detect the change with a minimal delay (ADD) while adhering to false alarm rate constraints (PFD). Consider a collection of obtained match scores, $x_1, x_2, \cdots, x_n$, from the AA system shown in Figure 2. Assuming that individual scores are mutually independent, QCD theory can be used to determine whether a change has occurred before time n or not. In the following subsections we present two main formulations of QCD.

A. Bayesian QCD (BQCD)

In the Bayesian setting [50], it is assumed that the time τ when the change occurs is distributed according to a geometric distribution, Geometric(ρ). Here, the value of ρ is the probability of a change occurring (an intrusion in this context). Conditioned on the change point τ, observations obtained before and after the change follows two distinct distributions, f_0 and f_1. At each time n, based on $\pi_i = P\{\tau = i\}$ for all $i < n$, a decision is made as to whether a change has occurred or not. Based on this formulation, ADD and PFD can be redefined as

$$ADD(\tau) = E[(\tau - T)^+] = \sum_{n=0}^{\infty} \pi_n E_n[(\tau - T)^+] \quad (2)$$

$$PFD(\tau) = P[\tau < T] = \sum_{n=0}^{\infty} \pi_n P_n[\tau < T], \quad (3)$$

where, for a Geometric(ρ) distribution,

$$\pi_n = P\{\tau = n\} = (1-\rho)^{n-1}\rho$$

for $0 < \rho < 1$ and $n > 0$. Then the Bayesian QCD becomes an optimization problem where the requirement is to minimize ADD subjected to a constraint on PFD. If the class of stopping times adhering to the constrain α on PFD is defined as

$$C_\alpha = \{\tau : PFD(\tau) < \alpha\},$$

then the QCD problem takes the form of Shiryaev's formulation [53], [50]. Objective of the Bayesian QCD formulated by Shiryaev is to obtain a stopping time $\tau \in C_\alpha$ to minimize ADD(τ) for a given α. If p_n is the posterior probability that a change has occurred at time n given observations up to time n

$$p_n = P[T \leqslant n | X^n],$$

where $X^n = (x_1, x_2, ..., x_n)$, then using the Bayes rule, it was shown in [50] that p_n follows a recursive formula as follows

$$p_{n+1} = \Phi(x_{n+1}, p_n),$$

where

$$\Phi(x_{n+1}, p_n) = \frac{\tilde{p_n} L(x_{n+1})}{\tilde{p_n} L(x_{n+1}) + (1 - \tilde{p_n})}.$$

Here, $\tilde{p_n} = p_n + (1 - p_n)\rho$ and

$$L(x_{n+1}) = \frac{f_1(x_{n+1})}{f_0(x_{n+1})}$$

is the likelihood ratio with $p_0 = 0$.

From Theorem 3.1 in [52], this recursive formula provides an optimal solution for the problem in hand with a stopping time of

$$\tau_s = \inf\{n \geqslant 1 : p_n \geqslant A\alpha\}$$

if $A \in (0, 1)$ can be chosen such that $PFD(\tau_s) = \alpha$. This method is known as the Shiryaev test and its proof can be found in [53], [52].

B. MiniMax QCD (MQCD)

In most of the practical AA systems, probability of intrusion is not known in advance. Therefore, it is important to study QCD in a non-Bayesian setting. MiniMax QCD formulation treats the change point τ as an unknown deterministic quantity [51], [52]. However, as earlier, it is assumed that pre-change distribution, f_0, and post-change distribution, f_1, are known. Due to the absence of prior knowledge on the change point, a reasonable measure of PFD is the reciprocal of mean time to a false detection as follows

$$PFD(\tau) = \frac{1}{E_\infty[\tau]}.$$

Based on this definition of PFD, Lorden proposed a minimax formulation for QCD [54], [51]. Consider the set of stopping times D_α for a given constraint α such that

$$D_\alpha = \{\tau : PFD(\tau) \leqslant \alpha\}.$$

Adhering to this constraint, Lorden's formulation optimizes a cost function to solve the minimax QCD problem. In particular, the cost function is the supremum of the average delay conditioned on the worst possible realizations as follows

$$WADD(\tau) = \sup_{n \geqslant 1} \operatorname{ess\,sup} E_n[(\tau - n)^+ | X^n].$$

Lorden's formulation tries to minimize the worst possible detection delay subjected to a constraint on PFD [54]. It was shown in [52], that the exact optimal solution for Lorden's formulation of QCD can be obtained using the CumSum algorithm [55].

C. CumSum Algorithm

Define the statistic $W(n)$ such that

$$W(n) = \max_{1 \leqslant k \leqslant n+1} \sum_{i=k}^{n} \log(L(x_i)),$$

and $W_0 = 0$, where $L(X_n) = f_1(X_n)/f_0(X_n)$ is the log likelihood ratio. It can be shown that the statistic $W(n)$ has the following recursive form

$$W_{n+1} = (W_n + \log(L(X_{n+1}))^+).$$

Time at which a change occured (τ) is chosen such that

$$\tau_c = \inf\{n \geqslant 1 : W_n \geqslant b\},$$

where b is a threshold. More details about the CumSum algorithm can be found in [55], [51], [52], [50].

D. Proposed Algorithm

Based on the Bayesian and MiniMax QCD algorithms, we propose an authentication algorithm to detect intrusions in an AA system. Essentially, our proposal is independent of all other base elements of an AA system (Figure 3). Therefore, existing AA systems can easily be extended to incorporate the proposed QCD method.

Training: In the training phase, the user is asked to perform a wide variety of tasks and sensor data are obtained. Predetermined features are then evaluated from the obtained sensor data. Part of the obtained features are stored in memory to serve as the gallery in the AA system. The remaining features are compared against chosen gallery to build a match distribution. In addition, the gallery entries are used to construct a non-match distribution based on the non-user features as illustrated in Figure 2. For the experiments conducted in this paper, a sample of other class data was used to model the non-match distribution. In practice, a common set of pre-obtained sensor data specific for the device can be used for this purpose. For example, face images of different users obtained from the same device can be made available in a cloud storage system for training.

input : Detection score of most recent iteration $score$, match score x_n, match distribution f_0, non-match distribution f_1, $Threshold$, $FloorThreshol$
output: Detection of an intrusion (Boolean)

//If it's the initial iteration set score to be zero;
if *isempty(score)* **then**
  $score = 0$;
else
  score = UpdateScore($score, x_n, f_0, f_1, FloorThreshold$);
  //FloorThreshold is used only resource efficient versions;
end
if $score > Threshold$ **then**
  Detect = True;
else
  Detect = False;
end
Return (Detect);

Algorithm 1: Main procedure proposed for decision making.

Testing: The proposed testing phase takes in to consideration a sequence of past observations when making a decision. At time n, the same set of sensor data and corresponding features g_n of the probe is collected as in the enrollment phase. Obtained features are compared against the signatures to obtain a score value x_n. A decision is made based on scores corresponding to all past observations $x_1, x_2, \cdots, x_n$ and the match distribution f_0 and non-match distribution f_1.

Described in Algorithm 1 is the proposed structure for decision making. A variable $score$ is initialized at zero and is updated using the method $UpdateScore$ once a new observation is observed. Once the score exceeds threshold A, a detection of a change is declared. In this paper we present two variants of the method $UpdateScore$ based on BQCD and MQCD. Those methods are listed in Algorithm 2 and Algorithm 3, respectively.

input : $score, x_n, f_0, f_1$
output: $score$

//Calculate likelihood ratio;
$L = f_1(x_n)/f_0(x_n)$;
$\tilde{p_n} = score + (1 - score)\rho$;
$score \leftarrow \frac{\tilde{p_n} L}{\tilde{p_n} L + (1 - \tilde{p_n})}$;
Return ($score$);

Algorithm 2: UpdateScore Method incorporating BQCD.

Illustrated in Figure 5 is the variation of detection scores when Bayesian QCD is used for the video shown in Figure 1. Detection scores values increase when there is significant variation in the expression. However, they decrease again once the neutral expression is returned. Since the intrusion occurs in Frame 201, the score value is seemed to be monotonically increasing. In this specific example, the likelihood ratio becomes infinity after the change point. Therefore, according to

input : $score, x_n, f_0, f_1$
output: $score$

//Calculate likelihood ratio;
$L = \log(f_1(x_n)/f_0(x_n))$;
$score \leftarrow score + L$;
Return ($score$);

Algorithm 3: UpdateScore Method incorporating MQCD.

Algorithm 2, the score is increasing by the assigned constant C. It should be noted that, slope of the curve could be increased by selecting a higher value for C in Algorithm 2. By the time the score passes the predetermined threshold, it is declared that an intrusion has occurred. For the set threshold in Figure 5, detection occurs with a delay of 9 samples.

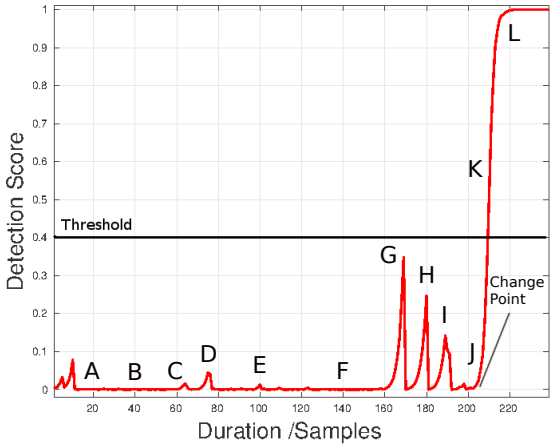

Fig. 5: Variation of Bayesian QCD scores for the video shown in Figure 1.

V. Resource Efficient Quickest Change Detection

In this section, we discuss how QCD can be performed with having a lower burden on the device resources. As noted in the introduction, a trivial solution to the problem of resource limitation is to perform detection on a few selected samples of observations. However, quickest detection performance may degrade greatly depending on how the sampling is done. In what follows, we introduce a data driven sampling rule based on data efficient QCD [51], [52], [56].

Consider a sequence of time instances $t = 1, 2, \cdots, i$ in which the device operates. At each time $i, i > 0$, a decision is made whether to take or skip an observation at time $i+1$. Let M_i be the indicator random variable such that $M_i = 1$ if the score x_i is used for decision making, and $M_i = 0$ otherwise. Thus, M_{i+1} is a function of the information available at time i, i.e. $M_{i+1} = \phi_i(I_i)$, where ϕ_i is the control law at time i, and $I_i = [M_1, M_2, \cdots, M_i, x_1^{M_1}, x_2^{M_2}, \cdots, s_i^{M_i}]$ represents the information at time i. Here, $x_i^{M_i}$ represents x_i if $M_i = 1$, otherwise x_i is absent from the information vector I_i. Let T be the stopping time on the information sequence $\{I_i\}$. Then, average percentage of observations (APO) obtained prior to the change point can be quantified as

$$APO = E\left[\frac{1}{T}\sum_{n=1}^{T} M_n\right]. \tag{4}$$

In the QCD scheme introduced in the previous section, observations are obtained at every time instant. Therefore, APO is equal to 1. A lower APO can be obtained while maintaining a lower ADD and PFD rates by employing an intelligent sampling mechanism. When such a mechanism is used, average usage of resources (memory, processing power, battery usage) are expected to decrease compared to the QCD scheme [47]. We introduce a technique to achieve this based on data efficient QCD.

A. Efficient Bayesian Formulation (E-BQCD)

In the Bayesian formulation of efficient QCD, an additional constraint based on the number of observations used is introduced in the optimization procedure. Define the Average Number of Observations (ANO) as

$$ANO = E\left[\sum_{n=1}^{min(\tau, T-1)} M_n\right]. \tag{5}$$

This quantity essentially captures the number of observations taken prior to the change point. It should be noted that ANO does not penalize additional observations taken after the change point. Therefore, ANO is a more conservative measure of the number of observations compared to APO, where $T \times APO \geq ANO$.

The efficient Bayesian QCD problem can be formulated as an optimization problem as follows[56],[52]

$$\begin{aligned} \underset{\phi,\tau}{\text{minimize}} \quad & ADD(\phi, \tau) \\ \text{subject to} \quad & PFA(\phi, \tau) \leq \alpha \\ & ANO(\phi, \tau) \leq \beta. \end{aligned} \tag{6}$$

In [56], an algorithm is presented to seek a possible solution for this optimization problem. Consider P_n, the probability that change had occurred by time n,

$$p_n = P(T \leq n | I_n),$$

where $p_0 = 0$. For $A, B \geq 0$ and $A > B$ the following control rule is proposed

$$M_{n+1} = \begin{cases} 0, & \text{if } p_n < B \\ 1, & \text{if } p_n \geq B. \end{cases}$$

Based on the value of M_{n+1}, p_{n+1} is updated as

$$p_{n+1} = \begin{cases} \tilde{p_n}, & \text{if } M_{n+1} = 0 \\ \frac{\tilde{p_n} L(x_{n+1})}{\tilde{p_n} L(x_{n+1}) + (1 - \tilde{p_n})}, & \text{if } M_{n+1} = 1, \end{cases}$$

where $\tilde{p_n} = p_n + (1 - p_n)\rho$ and $L(x_{n+1}) = \frac{f_1(x_{n+1})}{f_0(x_{n+1})}$. An intruder detection is declared at the earliest time (τ_D) when p_n surpasses the threshold A, i.e. $\tau_D = inf\{n \geq 1 : p_n > A\}$. It was proved in [56] that this algorithm is asymptotically optimal for the optimization formulation (6) for each fixed β when $\alpha \to 0$.

B. Efficient MiniMax Formulation (E-MQCD)

In a non-Bayesian setting, due to the absence of a priori distribution on the change point, a different quantity should be used to quantify the number of observations used for decision making. Work in [51],[52], proposes change Duty Cycle (CDC) as

$$CDC = \limsup_{n} \frac{1}{n} E_n \left[\sum_{k=1}^{n-1} M_k | \tau \geq n \right] \tag{7}$$

for this purpose. It should be noted that both CDC and APO are similar quantities. With the definition of CDC, efficient QCD in a minimax setting can be formulated as the following optimization problem

$$\begin{aligned} \underset{\phi,\tau}{\text{minimize}} \quad & ADD(\phi,\tau) \\ \text{subject to} \quad & PFA(\phi,\tau) \leq \alpha \\ & CDC(\phi,\tau) \leq \beta. \end{aligned} \tag{8}$$

In [51], a two threshold algorithm called DE-CumSum algorithm, is presented as a solution to this optimization problem. For suitably selected thresholds chosen to meet constraints α and β, it is shown to obtain the optimal lower bound asymptotically as $\alpha \to 0$. The DE-CumSum algorithm is presented below.

Start with $W_0 = 0$ and let $\mu > 0, A > 0$ and $h \geq 0$. For $n \geq 0$ use the following control rule

$$M_{n+1} = \begin{cases} 0 & \text{if } W_n < 0 \\ 1 & \text{if } W_n \geq 0. \end{cases}$$

Statistic W_n is updated as follows

$$W_{n+1} = \begin{cases} min(W_n + \mu, o), & \text{if } M_{n+1} = 0 \\ max(W_n + logL(X_{n+1}), -h), & \text{if } M_{n+1} = 1, \end{cases}$$

where $L(x) = \frac{f_1(x)}{f_0(x)}$. A change is declared at time τ_W, when the statistic W_n passes the threshold A for the first time as

$$\tau_W = inf\{n \geq 1 : W_n > A\}. \tag{9}$$

C. Modified Algorithm

Testing and training procedure under the resource efficient QCD-based detection is the same as proposed in Section IV-D. Testing is done using the $main$ method described in Algorithm 1 in section IV-D. Here, we present two alternative variants of the $UpdateScore$ method based on resource efficient BQCD and MQCD. Different steps are summarized in Algorithm 4 and Algorithm 5, respectively corresponding to the updates of E-BQCD and E-MQCD.

In Algorithm 8, parameter D is a constant. In our tests, this parameter was set to be equal to 1.Parameter $FloorThreshold$ is set equal to 0.05 in both algorithms. Evolution of score values when efficient Bayesian QCD is used is illustrated in Figure 6 for the case shown in Figure 1. In order to demonstrate the effect of using different sampling rates, the same experiment was conducted for a series of APO values. Functionality of efficient QCD algorithm can be explained using Figure 6. Consider the black line (corresponding to APO

input : $score, x_n, f_0, f_1, FloorThreshold$
output: $score$

//Calculate the priori probability $L = f_1(x_n)/f_0(x_n)$;
$\tilde{p_n} = score + (1 - score)\rho$;
//Use priori to update score when score is small **if** $score < FloorThreshold$ **then**
  $score = \tilde{p_n}$;
else
  $score \leftarrow \frac{\tilde{p_n} L}{\tilde{p_n} L + (1 - \tilde{p_n})}$;
end
Return ($score$);

Algorithm 4: UpdateScore method incorporating E-BQCD

input : $score, x_n, f_0, f_1, FloorThreshold$
output: $score$

if $score < 0$ **then**
  $score$ =min(score+D,0) ;
else
  $score \leftarrow$
  $max(score + log(\frac{f_1(x_n)}{f_0(x_n)}), -FloorThreshold)$;
end
Return ($score$);

Algorithm 5: UpdateScore method incorporating E-MQCD

= 0.92%) in Figure 6. After the initial observation at $t = 1$, no observations are taken until the score passes 0.05 at $t = 52$. In this duration, score is updated using a priori probability. Hence, the score is having a constant slope in this interval. At $t = 52$, as the score passes 0.05, an observation is taken and the score is updated based on log-likelihood as outlined in Algorithm 4. This causes a discontinuity in the graph by shifting the value of score onto 0.0007. Since this value is lower than 0.05, no observation is taken at $t = 53$. This process is continued until the score value surpasses the $Threshold$ value when an intrusion is declared.

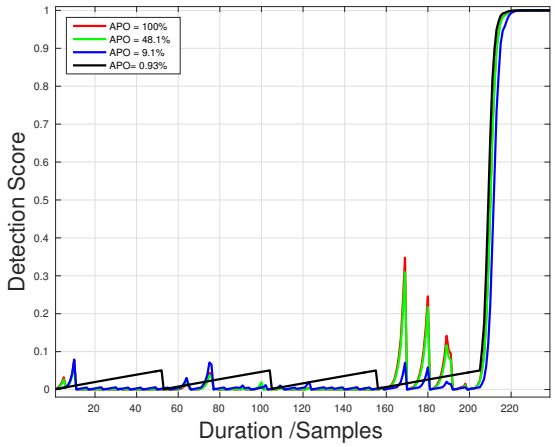

Fig. 6: Variations of efficient Bayesian QCD scores for the video shown in Figure 1 for different APO values.

Furthermore, Figure 6 suggests that variations of scores across time is somewhat similar when APO is 100% and 48.1% for the considered case. This shows that selecting

sampling points intelligently could reduce the sampling rate almost by half while producing near-identical performance for specific cases. The effect of sampling on the detection performance is discussed in detail in the following section.

VI. Experimental Results

A. Quickest Change Detection

We evaluated the performance of the proposed QCD methods using three publicly available unconstrained AA datasets - Touchalytics [17], MOBIO [11], and UMDAA-01 [12]. The following three previously proposed AA methods are used as the benchmark for comparisons.

Single score-based authentication (SSA): The present score value x_n alone is used to authenticate the user. If the score value is above a predetermined threshold, user is authenticated otherwise treated as an intrusion.

Time decay fusion (Sui et al.) [48]: In this method, two score samples fused by a linear function is used along with a decaying function to determine the authenticity of a user as, $s_n = wx_{n-1} + (1-w)x_n \times e^{\tau \delta t}$, where, w, τ are constants and δt is the time elapsed since the last observation.

Confidence functions (Crouse et al.) [15]: A sequential detection score S_{login} is calculated by incorporating time delay since the last observation and a function of the present score x_n. The detection score is evaluated as, $S_{login,n} = S_{login,n-1} + f_{map}(x_n) + \int_{t_{prev}}^{t_{now}} f_{dec} dt$. See [15] for the exact definitions of f_{map} and f_{dec}.

The PFD-ADD curves, introduced in Section III, are used to compare the performance of different methods. The PFD-ADD plot for the BQCD and MQCD methods can be obtained by varying the parameter $Threshold$ and plotting the ADD values corresponding to different PFD values. Similarly, the ADD-PFD curves for SSA and the method proposed by Sui et al. [48] and Crouse et al. [15] are obtained by varying the decision making threshold.

The measure of ADD signifies the latency of detecting an attack. On the other hand, PFD is a measure of false detections. A practical AA system should have a low latency in decision making as well as low false detection rate. Therefore, better AA systems are expected to have low ADD and PFD values. Hence, they should operate towards the lower left corner of the PFD-ADD curve, as illustrated in Figure 4. As a result, AA methods with very low operating values in the PFD-ADD plot are better in terms of their performance.

In the absence of a proper mobile dataset with intrusions, experimental data was obtained in the following manner for all datasets considered. For each dataset, all possible pairs of users were considered at a time. For each pair of users, full length signals (e.g. touch gestures or detected faces) of considered pair of users were merged to obtain a trial with a single intrusion. As a result, only one intruder/attacker was presented at each trial. Shown in Figure 1 is a sample trial obtained in this manner. Frames A to I correspond to the enrolled images of the genuine user. An intruder is presented at frame J and onwards. The intrusion point depends on the length of the samples corresponding to the first (genuine) user and therefore is not pre-determined. Each trial was tested using

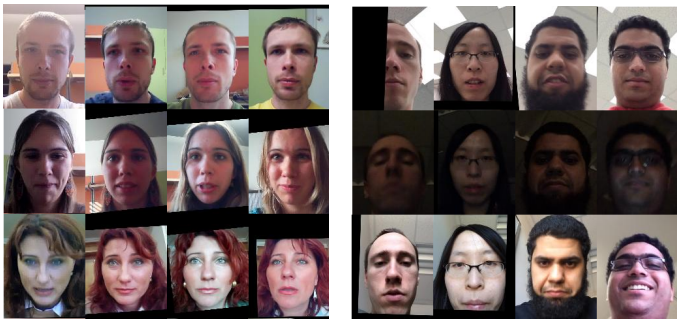

(a) (b)

Fig. 7: Sample detected face images from (a) the MOBIO dataset and (b) the UMDAA-01 dataset.

before mentioned methods to determine detection delay and probability of false detections under each method.

1) UMDAA-01 Dataset: The UMDAA-01 dataset [12] consists of images of 50 individuals taken from an iPhone 5 device across three sessions performing five tasks including an enrollment task. Both face images as well as touch gestures are simultaneously captured in this dataset. Sample detected face images from this dataset are shown in Figure 7(b). As suggested in [12], enrollment data was used as gallery and data from the other sessions was used as probes. In addition 20 number of instances from the probe session was used to obtain the match score distribution. When testing, 33 % of the remaining subjects excluding the probe class and the target class were randomly chosen to obtain the non-match distribution.

Results on the Face Data: Face images of the user were normalized and image regions corresponding to eyes, nose, lips and eyebrows were extracted. The HOG features [57] were extracted on each facial component. These features were concatenated to obtain the resulting feature for the given face. Cosine distance is used to generate score values by matching enrollment data with probes. Figure 8 shows the ADD-PFD plot corresponding the UMDAA-01 face data. From this figure, it can be seen that both BQCD (ρ = 0.001) and MQCD outperform the other methods. This can be seen by comparing their performances in the low PFD region.

Results on the Touch Data: From each swipe data, a 27-dimensional feature vector is extracted using the method described in [17]. A single class SVM with RBF kernel was used to generate matching scores. Figure 9 shows the ADD-PFD curves corresponding to different methods on this dataset. It should be noted that there exists a considerable similarity between single touch swipes of different users. Therefore, from Figure 9, methods that rely on data of single or two swipes have performed poorly. It can be seen that BQCD, MQCD and the method proposed by Crouse et al. [15] that uses information from pre and post change distributions have performed reasonably well. In general, the MQCD method yields faster detection rates and low false detections compared to the other methods.

2) MOBIO Dataset: The MOBIO dataset [11] contains videos of 152 subjects taken across two phases where each phase consists of six sessions each. Videos in this dataset are

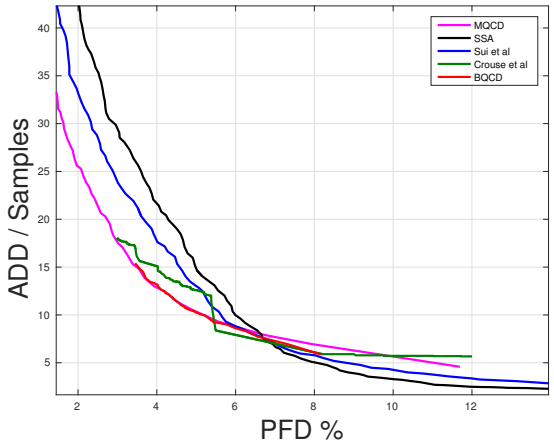

Fig. 8: Performance curves obtained on the UMDAA-01 face dataset.

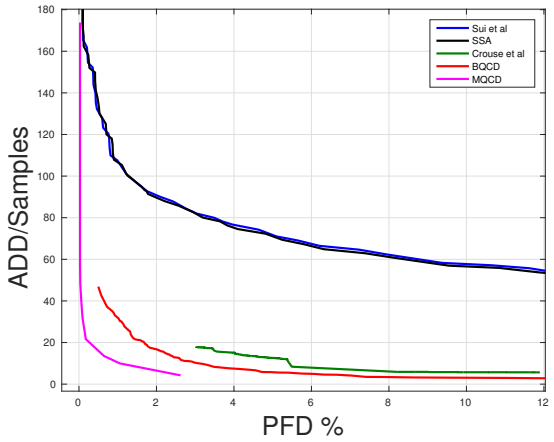

Fig. 9: Performance curves obtained on the UMDAA-01 touch dataset.

acquired using a standard 2008 Macbook laptop computer and a NOKIA N93i mobile phone (See Figure 7(a)). Following the protocol defined in [13], video frames of the 12th session were considered as the enrollment data and video frames of all other sessions were used as probes. We conducted our experiments on the laptop image data based on the LBP features. Again, the cosine distance was used to generate the match and non-match scores. Figure 10 shows the performance curves corresponding to different methods on the MOBIO dataset. Note that the images in this dataset are well aligned and mostly frontal. As a result, pre-change and post-change distributions are well separated. Hence, all considered methods yielded relatively better performance. However, the BQCD and MQCD methods have performed marginally better than the other compared approaches.

3) Touchalytics Dataset: The Touchalitics dataset contains touch data of 37 users collected across 7 tasks. Similar to the UMDAA-01 touch dataset, touch gesture features are extracted using the method described in [17] and a single class SVM with RBF kernel was used to generate match and non-match scores. Figure 11 shows the performance of different methods on this dataset. As before, making a decision based on a single

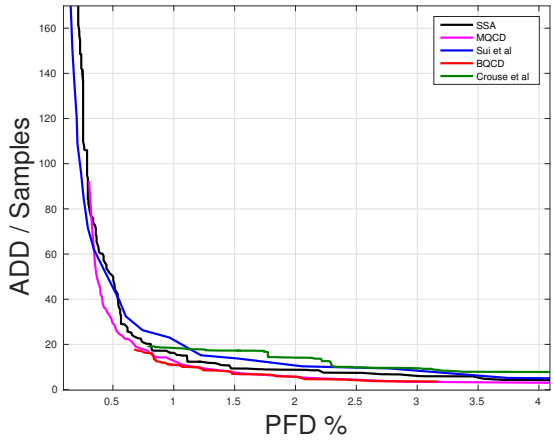

Fig. 10: Performance curves obtained on the MOBIO face dataset.

swipe or two swipes have appeared to perform poorly. The MQCD method performs the best followed by the BQCD method and the method of Crouse et al. [15].

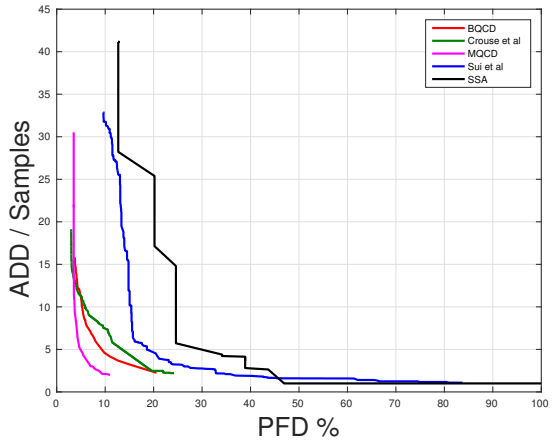

Fig. 11: Performance curves obtained on Touchalytics dataset.

4) Discussion: From the above experiments, it can be seen that the BQCD and MQCD methods have outperformed the other existing AA methods. Furthermore, in all cases, the MQCD method has performed marginally better than the BQCD method. This is mainly due to the error induced by approximating the change distribution by a Geometric (ρ) distribution. In practice, where information on change (intrusion) probability is unknown in advance, the MQCD method provides more usability as opposed to the BQCD method.

Detection delay and probability of false detections of the proposed algorithm depend on the type of features as well as the classifiers used for matching. The proposed method is not restricted to any specific type of feature or a classifier. Therefore, by using better features and classifiers it is possible to obtain even lower ADD and PFD values.

Furthermore, it should be noted that, the detection delay rates (ADD) shown in Figures 8, 9, 10, and 11 are highly inflated as a result of non-detected intrusions due to the limitations of the features and/or classifiers. To further elaborate

on this point, let us consider the implementation of the MQCD method with a threshold chosen such that PFD is at 5%. Tabulated in Table III is the distribution of detection delay (ADD) for the tests conducted. According to Table III, nearly 90% of the time, an intrusion can be detected using less than 7 samples. Therefore, the proposed method would produce quick results for a small false detection rate in a practical setting.

	2-3 S	4-5 S	6-7 S	8-10S	>10 S
UMD-Face	11.9	17.07	55.06	6.02	9.93
UMD-Touch*	73.62	13.51	4.69	3.04	3.13
MOBIO	8.74	61.87	10.38	7.51	11.5
Touchalytics	3.65	7.23	82.23	2.94	3.94
Mean	24.47	24.92	38.09	4.87	7.12

TABLE III: Percentage breakdown of delay times (in samples) for a fixed PFD of 5% for MQCD. *3% of PFD was used instead.

B. Resource Efficient QCD

Effect of extending QCD to incorporate resource efficiency through sampling was studied on the before mentioned three datasets. Performance of the proposed sampling method was compared against the following two benchmark sampling methods.

Fixed Time Step Sampling: Most of the existing AA systems employ a sampling mechanism where sensor observations are obtained with a fixed inter-sample interval [15],[48]. In our experiments, this interval was chosen to satisfy the given APO rate.

Dice Sampling: In this method, a weighted coin is tossed at every time instant to determine whether a sample should be obtained or not [56],[51],[52]. The weight of the coin is equal to the chosen APO value.

Same set of features and classifiers as described in Section VI-A were used to evaluate performance of the proposed methods. For each dataset considered in this paper, E-BQCD and E-MQCD were applied on top of BQCD and MQCD for a specific APO rate. In addition, BQCD method was implemented using time step sampling and DICE sampling for comparison.

Shown in Figure 12 are the performance curves obtained for the UMDAA-01 face dataset for an APO of 21%. Performance curves have shifted to the left by some margin and have moved slightly upwards as shown in the graph due to sampling. At a glance, performance appears to have improved despite lower sampling for a given PFD value. In comparison, sampling with DICE and fixed time step has worsen the initial result. The same trend seems to follow in the UMDAA-01 touch dataset as seen in Figure 13 for an APO of 17%. Although BQCD tends to perform poorly compared to MQCD, performance of resource efficient versions of BQCD and MQCD are comparable.

Results obtained for the experiments done on the MOBIO face dataset for an APO of 17% are shown in Figure 14. Both QCD methods yielded comparable results on the MOBIO

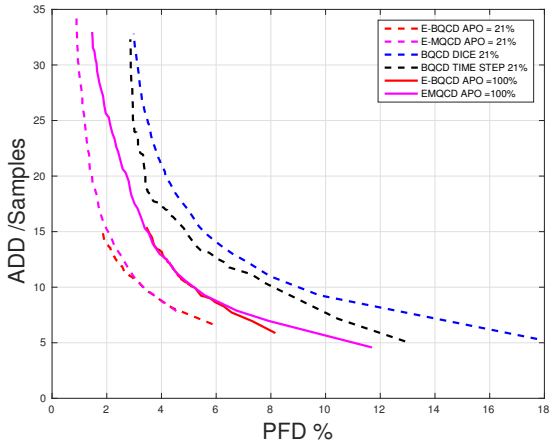

Fig. 12: Performance curves obtained on the UMDAA-01 face dataset for efficient QCD.

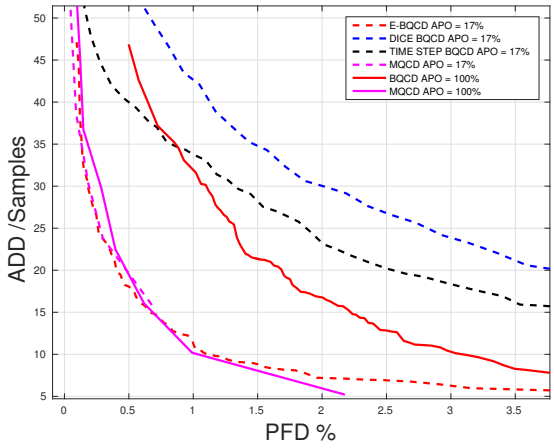

Fig. 13: Performance curves obtained on the UMDAA-01 touch dataset for efficient QCD.

dataset in our earlier experiments. When resource efficient QCD was employed, performance curves for both methods improved nearly by an equal amount compared to the QCD performance. It should be noted that, both DICE and fixed time sampling performances are much worse compared to E-QCD on the MOBIO dataset.

Final set of experiments were carried out on the Touchalytics touch dataset with an APO rate of 17%. Results of these experiments are presented in Figure 15. As in earlier cases, resource efficient QCD has outperformed QCD and other sampling methods. However, there are a couple of notable differences. Unlike in earlier experiments, E-BQCD and E-MQCD performance curves do not overlap in this case. However, this is only due to the absence of a common operating region. In addition, time step sampling performed better than DICE sampling on this dataset.

Resource efficient QCD have improved the performance of QCD and has performed better than alternative sampling methods have. The exact shape of the performance curves and gaps between each curves depend on the type of feature and classifier used. Irrespective of this, efficient QCD has yielded better results on average. This can be seen from the results

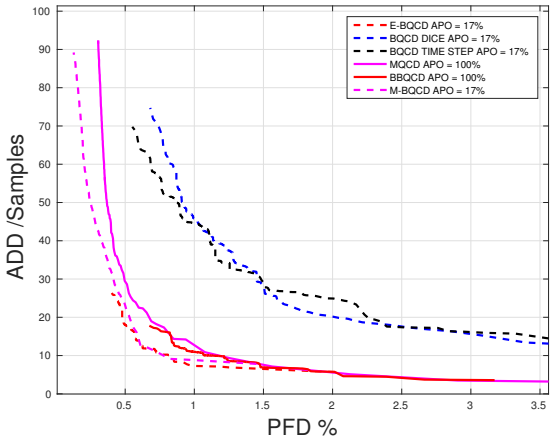

Fig. 14: Performance curves obtained on the MOBIO face dataset for efficient QCD.

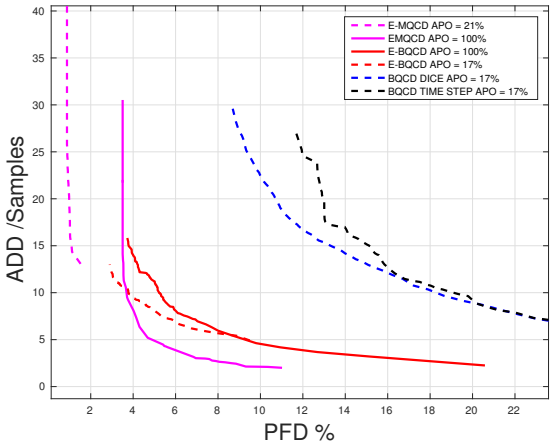

Fig. 15: Performance curves obtained on the Touchalytics touch dataset for efficient QCD.

summarized in Table IV, where PFD values obtained for a fixed detection delay of 15 samples for all considered datasets are listed.

In order to investigate this phenomena further, we carried out a case study on the UMDAA-01 face dataset. We conducted the above mentioned experiment on the dataset for a range of APO values. The resulted performance curves are shown in Figure 16. Performance curves in Figure 16 suggest that as APO decreases, performance curves keep on shifting further left. However, at the same time, the minimum possible detection time has also increased. Therefore, very low sampling is not feasible if quick change detection is desired. On the other hand, for a fixed detection delay, it might be possible to select a lower sampling rate so that lower PFD is obtained. This result is true for all datasets we considered as evident from Table IV.

In Figure 17, we plot the minimum possible detection time for different APO values for the test conducted on UMDAA-01 face dataset. As evident from this figure, the minimum detection times increase as sampling fraction (APO) is increased. Therefore, for practical applications, it is desired to select a moderate value for APO when efficient QCD is used.

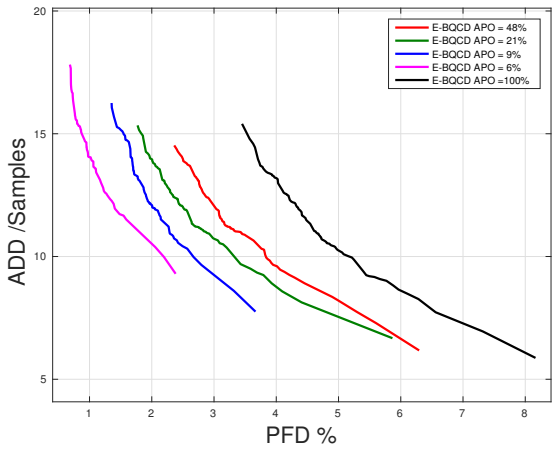

Fig. 16: Effect of using different APO values for sampling.

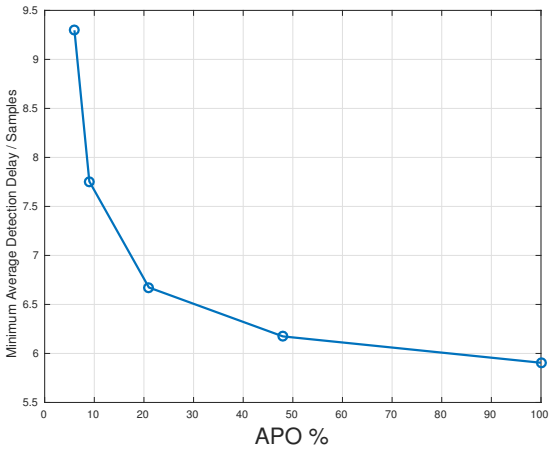

Fig. 17: Effect of sampling on minimum average detection delay.

1) Discussion: Decreasing the sampling rate by selecting a threshold to achieve a lower APO have decreased the PFD while increasing the ADD for all the cases considered. This observation can be justified based on the score update mechanism. Consider the black curve (APO = 0.93%) and the blue curve (APO = 9.1 %) in the score evolution shown in Figure 6. Note that, with the way how sampling is performed, the black curve had missed all the humps created due to irregularities in sensing. The blue curve on the other hand, is affected only by a few humps due to sampling. However, due to low sampling in humps, the blue curve has not risen as much as the red curve (APO = 100%) has. This suggests that for a constant threshold, occurrences of false detections will be lower for the blue curve (APO = 9.1 %). Therefore, PFD decreases when sampling is carried out. On the other hand, due to sampling, an intrusion may not be sensed till the sampling that follows is carried out. This is clearly seen in the case of the blue line in Figure 6. As a result, detection delay increases for a fixed threshold when sampling is carried out as shown in Figure 17.

	SSA	Sui et al [48]	Crouse et al [15]	BQCD	MQCD	E-BQCD	E-MQCD	DICE [52]	TIMESTEP [48],[15]
UMD-Face*	5.1	4.6	4.0	3.5	3.4	**1.8**	2.1	5.5	4.8
UMD-Touch	48.8	50.2	4.0	2.2	**0.6**	**0.6**	**0.6**	5.4	3.9
MOBIO	1.1	1.2	1.7	0.8	0.8	**0.5**	0.6	3.0	3.5
Touchalytics	24.6	14.8	3.1	3.8	3.5	NA	**1.0**	13.1	15.0

TABLE IV: The PFD % rate for a detection delay of 15 samples. The efficient QCD methods use an APO of 17%. * APO of 21% used.

VII. Conclusion

In this paper we addressed the issue of detecting intrusions in a mobile AA system with quickest time with high utilization of resources. We presented a method for detecting an intrusion in an AA system with a minimal delay with a constraint on false detection rate. Two variants of the QCD based on Bayesian and MiniMax formulations were introduced. Performance of the proposed method was demonstrated using three publicly available datasets.

The basic QCD methodologies were extended using resource efficient QCD where a data driven observation sampling was introduced with the aim of increasing resource efficiency. The introduced algorithms not only reduced number of observations taken, but also improve the performance of the system in terms of latency and false detections. Validity of this result was demonstrated using various datasets.

The proposed method does not rely on a specific feature or a classifier for its performance. This was verified in testing by using different classifiers and features for different datasets. Therefore, existing AA methods can be extended using the proposed method to enhance the performance. It was shown that the proposed method is effective even when there is a considerable overlap between pre and post-change distributions.

Acknowledgement

This work was supported by US Office of Naval Research (ONR) Grant YIP N00014-16-1-3134.

References

[1] lookout.com. (2015) Phone theft in america. [Online]. Available: http://www.lookout.com/resources/reports/phone-theft-in-america

[2] C. Deitrick. (2015) Smartphone thefts drop as kill switch usage grows. [Online]. Available: http://http://www.consumerreports.org/cro/news/2015/06/smartphone-thefts-on-the-decline/index.htm

[3] W. Meng, D. S. Wong, S. Furnell, and J. Zhou, "Surveying the development of biometric user authentication on mobile phones," *IEEE Communications Surveys Tutorials*, vol. 17, no. 3, pp. 1268–1293, 2015.

[4] D. Tapellini. (2014, May) Smart phone thefts rose to 3.1 million in 2013: Industry solution falls short, while legislative efforts to curb theft continue. [Online; posted May 28, 2014]. [Online]. Available: http://www.consumerreports.org/cro/news/2014/04/smart-phone-thefts-rose-to-3-1-million-last-year/index.htm

[5] H. Khan, U. Hengartner, and D. Vogel, "Usability and security perceptions of implicit authentication: Convenient, secure, sometimes annoying," in *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, 2015, pp. 225–239.

[6] S. Egelman, S. Jain, R. S. Portnoff, K. Liao, S. Consolvo, and D. Wagner, "Are you ready to lock?" in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014, pp. 750–761.

[7] M. Harbach, E. von Zezschwitz, A. Fichtner, A. D. Luca, and M. Smith, "It's a hard lock life: A field study of smartphone (un)locking behavior and risk perception," in *Symposium On Usable Privacy and Security (SOUPS 2014)*, 2014, pp. 213–230.

[8] I. T. Fischer, C. Kuo, L. Huang, and M. Frank, "Short paper: Smartphones: Not smart enough?" in *ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, ser. SPSM '12. New York, NY, USA: ACM, 2012, pp. 27–32. [Online]. Available: http://doi.acm.org/10.1145/2381934.2381941

[9] A. J. Aviv, K. Gibson, E. Mossop, M. Blaze, and J. M. Smith, "Smudge attacks on smartphone touch screens," in *Proceedings of the 4th USENIX Conference on Offensive Technologies*, 2010, pp. 1–7.

[10] G. D. Clark and J. Lindqvist, "Engineering gesture based authentication systems," *IEEE pervasive computing mobile and ubiquitous systems*, vol. 14, no. 1, pp. 18–25, 2015.

[11] C. McCool, S. Marcel, A. Hadid, M. Pietikainen, P. Matejka, J. Cernocky, N. Poh, J. Kittler, A. Larcher, C. Levy, D. Matrouf, J.-F. Bonastre, P. Tresadern, and T. Cootes, "Bi-modal person recognition on a mobile phone: Using mobile phone data," in *IEEE International Conference on Multimedia and Expo Workshops*, July 2012, pp. 635–640.

[12] M. E. Fathy, V. M. Patel, and R. Chellappa, "Face-based active authentication on mobile devices," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2015.

[13] P. Samangouei, V. M. Patel, and R. Chellappa, "Attribute-based continuous user authentication on mobile devices," in *IEEE International Conference on Biometrics: Theory, Applications and Systems*, 2015.

[14] A. Hadid, J. Heikkila, O. Silven, and M. Pietikainen, "Face and eye detection for person authentication in mobile phones," in *ACM/IEEE International Conference on Distributed Smart Cameras*, Sept 2007, pp. 101–108.

[15] D. Crouse, H. Han, D. Chandra, B. Barbello, and A. K. Jain, "Continuous authentication of mobile user: Fusion of face image and inertial measurement unit data," in *International Conference on Biometrics*, 2015.

[16] P. Perera and V. M. Patel, "Towards multiple user active authentication in mobile devices," in *IEEE International Conference on Automatic Face and Gesture Recognition*, 2017.

[17] M. Frank, R. Biedert, E. Ma, I. Martinovic, and D. Song, "Touchalytics: On the applicability of touchscreen input as a behavioral biometric for continuous authentication," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 1, pp. 136–148, Jan 2013.

[18] A. Serwadda, V. Phoha, and Z. Wang, "Which verifiers work?: A benchmark evaluation of touch-based authentication algorithms," in *IEEE International Conference on Biometrics: Theory, Applications and Systems*, Sept 2013, pp. 1–8.

[19] T. Feng, Z. Liu, K.-A. Kwon, W. Shi, B. Carbunar, Y. Jiang, and N. Nguyen, "Continuous mobile authentication using touchscreen gestures," in *IEEE Conference on Technologies for Homeland Security*, Nov 2012, pp. 451–456.

[20] M. Sherman, G. Clark, Y. Yang, S. Sugrim, A. Modig, J. Lindqvist, A. Oulasvirta, and T. Roos, "User-generated free-form gestures for authentication: Security and memorability," in *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*, 2014, pp. 176–189.

[21] X. Zhao, T. Feng, W. Shi, and I. Kakadiaris, "Mobile user authentication using statistical touch dynamics images," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 11, pp. 1780–1789, Nov 2014.

[22] H. Zhang, V. M. Patel, M. E. Fathy, and R. Chellappa, "Touch gesture-based active user authentication using dictionaries," in *IEEE Winter conference on Applications of Computer Vision*, 2015.

[23] C. Bo, L. Zhang, X.-Y. Li, Q. Huang, and Y. Wang, "Silentsense: Silent user identification via touch and movement behavioral biometrics," in

Proceedings of the 19th Annual International Conference on Mobile Computing & Networking, ser. MobiCom '13. New York, NY, USA: ACM, 2013, pp. 187–190.

[24] L. Fridman, S. Weber, R. Greenstadt, and M. Kam, "Active authentication on mobile devices via stylometry, gps location, web browsing behavior, and application usage patterns," *IEEE Systems Journal*, 2015.

[25] J. Kwapisz, G. Weiss, and S. Moore, "Cell phone-based biometric identification," in *IEEE International Conference on Biometrics: Theory Applications and Systems*, Sept 2010, pp. 1–7.

[26] M. Derawi, C. Nickel, P. Bours, and C. Busch, "Unobtrusive user-authentication on mobile phones using biometric gait recognition," in *International Conference on Intelligent Information Hiding and Multimedia Signal Processing*, Oct 2010, pp. 306–311.

[27] A. Primo, V. Phoha, R. Kumar, and A. Serwadda, "Context-aware active authentication using smartphone accelerometer measurements," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, June 2014, pp. 98–105.

[28] H. Zhang, V. M. Patel, S. Shekhar, and R. Chellappa, "Domain adaptive sparse representation-based classification," in *IEEE International Conference on Automatic Face and Gesture Recognition*, vol. 1, May 2015, pp. 1–8.

[29] P. Perera and V. M. Patel, "Extreme value analysis for mobile active user authentication," in *IEEE International Conference on Automatic Face and Gesture Recognition*, 2017.

[30] M. Abramson and D. W. Aha, "User authentication from web browsing behavior," in *Florida Artificial Intelligence Research Society Conference*, AAAI Press. AAAI Press, 2013.

[31] F. Li, N. Clarke, M. Papadaki, and P. Dowland, "Active authentication for mobile devices utilising behaviour profiling," *International Journal of Information Security*, vol. 13, no. 3, pp. 229–244, 2014.

[32] H. Saevanee, N. Clarke, S. Furnell, and V. Biscione, "Text-based active authentication for mobile devices," in *ICT Systems Security and Privacy Protection*, ser. IFIP Advances in Information and Communication Technology, N. Cuppens-Boulahia, F. Cuppens, S. Jajodia, A. Abou El Kalam, and T. Sans, Eds. Springer Berlin Heidelberg, 2014, vol. 428, pp. 99–112.

[33] H. Gascon, S. Uellenbeck, C. Wolf, and K. Rieck, "Continuous authentication on mobile devices by analysis of typing motion behavior," in *Sicherheit 2014*, 2014, pp. 1–12.

[34] C. Holz, S. Buthpitiya, and M. Knaust, "Bodyprint: Biometric user identification on mobile devices using the capacitive touchscreen to scan body parts," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2015, pp. 3011–3014.

[35] V. M. Patel, R. Chellappa, D. Chandra, and B. Barbello, "Continuous user authentication on mobile devices: Recent progress and remaining challenges," *IEEE Signal Processing Magazine*, vol. 33, no. 4, pp. 49–61, July 2016.

[36] H. Zhang, V. M. Patel, and R. Chellappa, "Robust multimodal recognition via multitask multivariate low-rank representations," in *IEEE International Conference on Automatic Face and Gesture Recognition*, vol. 1, May 2015, pp. 1–8.

[37] K. B. Schaffer, "Expanding continuous authentication with mobile devices," *Computer*, vol. 48, no. 11, pp. 92–95, Nov 2015.

[38] R. P. Guidorizzi, "Security: Active authentication," *IT Professional*, vol. 15, no. 4, pp. 4–7, July 2013.

[39] M. Jakobsson, E. Shi, P. Golle, and R. Chow, "Implicit authentication for mobile devices," in *Proceedings of USENIX*, 2009.

[40] E. Shi, Y. Niu, M. Jakobsson, and R. Chow, "Implicit authentication through learning user behavior," in *Proceedings of the 13th International Conference on Information Security*, 2011, pp. 99–113.

[41] N. L. Clarke, *Transparent User Authentication - Biometrics, RFID and Behavioural Profiling*. Springer, 2011.

[42] N. Clarke, S. Karatzouni, and S. Furnell, *Emerging Challenges for Security, Privacy and Trust: 24th IFIP TC 11 International Information Security Conference, SEC 2009, Pafos, Cyprus, May 18–20, 2009. Proceedings*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, ch. Flexible and Transparent User Authentication for Mobile Devices, pp. 1–12.

[43] H. Crawford and K. Renaud, "Understanding user perceptions of transparent authentication on a mobile device," *Journal of Trust Management*, vol. 1, no. 7, pp. 1–28, 2014.

[44] J. G. W. Ding Li, Shuai Hao and G. J. Halfond, "An empirical study of the energy consumption of android applications," in *IEEE International Conference on Software Maintenance and Evolution (ICSME)*, September 2014.

[45] W. Lee. (2013) Mobile apps and power consumption - basics, part 1. [Online]. Available: https://developer.qualcomm.com/blog/mobile-apps-and-power-consumption-basics-part-1

[46] P. Samangouei, V. M. Patel, and R. Chellappa, "Facial attributes for active authentication on mobile devices," *Image and Vision Computing*, May 2016.

[47] P. Perera and V. M. Patel, "Quickest intrusion detection in mobile active user authentication," in *International Conference on Biometrics Theory, Applications and Systems*, 2016.

[48] a. Z. Y. Sui, E.Y.Du, and F.Li, "Secure and privacy-preserving biometrics based active authentication," in *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2012, pp. 1291 – 1296.

[49] C. Braz and J.-M. Robert, "Security and usability: The case of the user authentication methods," in *Proceedings of the 18th Conference on L'Interaction Homme-Machine*, 2006, pp. 199–203.

[50] V. V. Veeravalli and T. Banerjee, "Quickest Change Detection," *ArXiv e-prints*, Oct. 2012.

[51] T. Banerjee and V. Veeravalli, "Data-efficient quickest change detection in minimax settings," *IEEE Transactions on Information Theory*, pp. 6917 – 6931, Oct 2013.

[52] T. Banerjee and V. V. Veeravalli, "Data-efficient quickest change detection," *Sri Lankan Journal of Applied Statistics, Special Issue: Modern Statistical Methodologies in the Cutting Edge of Science*, pp. 183–208, Nov 2014.

[53] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory of Probability & Its Applications*, vol. 8, no. 1, pp. 22–46, 1963.

[54] G. Lorden, "Procedures for reacting to a change in distribution," *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, 12 1971.

[55] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1/2, pp. 100–115, 1954.

[56] T. Banerjee and V. V. Veeravalli, "Data-Efficient Quickest Change Detection with On-Off Observation Control," *ArXiv e-prints*, May 2011.

[57] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, 2005, pp. 886–893.

Pramuditha Perera Pramuditha Perera is a Ph.D. candidate at the Department of Electrical and Computer Engineering, Rutgers University, New Jersey, USA. He received his bachelors degree in Electrical and Electronic Engineering from University of Peradeniya, Sri Lanka in 2014. His research interests include computer vision and machine learning with applications in biometrics.

Vishal M. Patel Vishal M. Patel [SM'16] is an A. Walter Tyson Assistant Professor in the Department of Electrical and Computer Engineering at Rutgers University. Prior to joining Rutgers University, he was a member of the research faculty at the University of Maryland Institute for Advanced Computer Studies (UMIACS). He completed his Ph.D. in Electrical Engineering from the University of Maryland, College Park, MD, in 2010. His current research interests include signal processing, computer vision, and pattern recognition with applications in biometrics and imaging. He has received a number of awards including the 2016 ONR Young Investigator Award, the 2016 Jimmy Lin Award for Invention, A. Walter Tyson Assistant Professorship Award, Best Paper Award at IEEE AVSS 2017, Best Paper Award at IEEE BTAS 2015, and Best Poster Awards at BTAS 2015 and 2016. He is an Associate Editor of the IEEE Signal Processing Magazine, IEEE Biometrics Compendium, and serves on the Information Forensics and Security Technical Committee of the IEEE Signal Processing Society. He is a member of Eta Kappa Nu, Pi Mu Epsilon, and Phi Beta Kappa.