Quickest Intrusion Detection in Mobile Active User Authentication

Pramuditha Perera and Vishal M. Patel
Department of Electrical and Computer Engineering
Rutgers University, CoRE 508, 94, Brett Road, Piscataway, NJ, 08854, USA
pramuditha.perera@rutgers.edu, vishal.m.patel@rutgers.edu *

Abstract

In this paper, we address the problem of quickly detecting intrusions with lower false detection rates in mobile Active Authentication (AA) systems. Bayesian and Minimax versions of the Quickest Change Detection (QCD) algorithms are introduced to quickly detect intrusions in mobile AA systems. Furthermore, we introduce a new evaluation metric for comparing the performance of different AA systems. Effectiveness of the proposed framework is demonstrated using three publicly available unconstrained AA datasets. It is shown that the proposed QCD-based intrusion detection method can perform better than many traditional AA methods in terms of latency and low false detection rates.

1. 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. Therefore, modern mobile devices hold a substantial intangible value in terms of content in addition to their physical value. In this context, information security in mobile devices has become a key concern for the consumers. According to [11], one third of smartphone users are likely to pay up to $1,000 to retrieve stolen phone data alone. This amount is a reflection of the value of the sensitive information modern mobile devices hold.

In order to overcome this issue of security and privacy of mobile devices, many smartphone device vendors have introduced a remotely operable kill switch that can wipe out data from the mobile memory [6]. However, in the case of a stolen device, there is no guarantee that information theft wouldn't occur by the time when the kill switch is operated. On the other hand, mobile devices needn't necessarily be stolen for information theft to take place. Information theft could take place when an intruder or attacker has physical access to the device; typically when the device is left behind in public places or during a burglary even without the attention of the owner [11]. Therefore, it is important to actively monitor the possibly of an intrusion to reduce the risk of information theft in mobile devices.

Figure 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.

To deal with this issue, Active Authentication (AA), where users' identity is continuously monitored and verified, has emerged as a promising solution for the information security problem in mobile devices [15]. Recently proposed AA systems rely on various sensor information to arrive at physiological and behavioral profiles of the genuine user to continuously verify his/her identity. For example, a user's face, which is a physiological biometric, can be captured using the front-facing camera of a mobile device and can be used to continuously authenticate a mobile device user [9], [13], [4], [7], [16]. On the other hand, sensors such as gyroscope, touchscreen and accelerometer can be used to measure behavioral biometric traits such as gait [23], [10], touch gestures [8], [17], [22] and hand movement [19] transparently. It should be noted that most of these works concentrate on producing features and classifiers that produce higher intrusion detection rates.

However, to the best of our knowledge, how fast an AA system could detect an attacker has not been widely studied in the literature. Yet, it remains to be an important feature of an AA system. For example, if an AA system requires 3-

*This work was supported by cooperative agreement FA8750-13-2-0279 from DARPA.

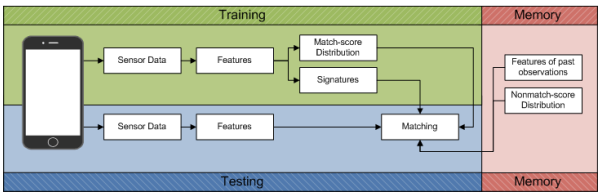

Figure 2. An overview of the proposed QCD-based AA method.

4 minutes to detect an intrusion successfully, it would grant an intruder plenty of time to extract required 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. For example, consider a series of observations captured from the front-facing camera of a mobile device as shown in Figure 1. Frames (A-I) in Figure 1 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 main objective of an AA system is to detect intrusions with a minimal delay while maintaining a low rate of false detections. For instance, note the changes in facial expression of the genuine user in frames (D-F). While having a fast response, an AA system ideally should not falsely interpret this expression change as an intrusion.

Most of the existing AA systems rely on a single observation to determine the authenticity of the user [16], [8], [7]. Intruder detection scheme proposed in [21] 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 [4] 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.

In this paper, 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 prevailing performance. Furthermore, we introduce a set of parameters that can be used to evaluate the performance of the intruder detection scheme of an AA system.

Remainder of the paper is organized as follows. Section 2 describes the quality measures for intruder detection in AA. In Section 3, we first briefly introduce the idea behind QCD and then present our proposed framework for QCD-based intrusion detection in mobile AA systems. Experimental results are presented in Section 4 and Section 5 concludes the paper with a brief summary and discussion.

2. 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 signatures are constructed using features of the enrolled data. These set of signatures serve as the gallery at the matching stage. When a probe is present, the same set of features as before are obtained from the sensor data and are compared against the gallery using a matching rule. 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 enter a password. Otherwise the user is allowed to continue with the device until the next sensor observation.

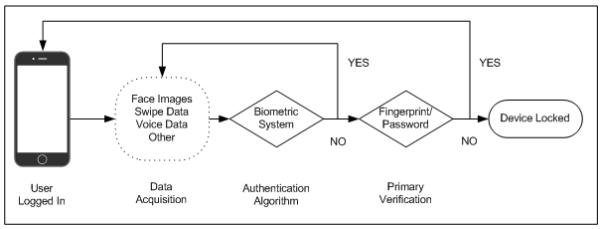

Figure 3. An overview of a typical AA system.

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 consid-

erably. 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.

2.1. 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 [3], it is more desirable to have an AA system with a low intrusion detection delay.

2.2. 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 [3] 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 is a typical ADD-PFD plot drawn for a non-sequential AA system. 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.

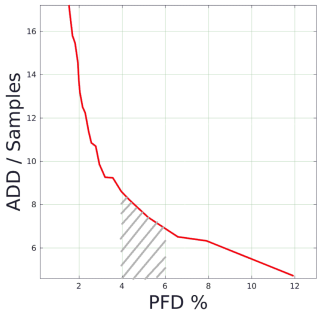

Figure 4. A typical PFD-ADD curve.

Adhering to security and usability principles [3], 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 systems should operate in a region where both ADD and PFD are lower. Based on this rationale, operable area of an AA system is the shaded area as shown in Figure 4.

3. 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 [20], [1], [2]. 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.

3.1. Bayesian QCD (BQCD)

In the Bayesian setting [20], 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 $\pi_i = P\{\tau = n\}$ for all $i < n$, a decision is made as to whether a change has occurred or not. Based on this formu-

lation, ADD and PFD can be redefined as

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

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

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 [18], [20]. 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 [20] 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 [2], 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 [18], [2].

3.2. 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 [1], [2]. 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 [12], [1]. 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 [12]. It was shown in [2], that the exact optimal solution for Lorden's formulation of QCD can be obtained using the CumSum algorithm [14].

3.3. 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 [14], [1], [2], [20].

3.4. 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 signatures to build a match distribution. In addition, the same set of signatures 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, parameter ρ, $Threshold$
output: Detection of an intrusion (Boolean)

//If it's the initial iteration set score to be zero;
if *isempty(*$score$*)* **then**
  $score = 0$;
else
end
//Calculate likelihood ratio;
$L = \log(f_1(x_n)/f_0(x_n))$;
$\tilde{p_n} = score + (1 - score)\rho$;
$score \leftarrow \frac{\tilde{p_n} L}{\tilde{p_n} L + (1 - \tilde{p_n})}$;
if $score > Threshold$ **then**
  Detect = True;
else
  Detect = False;
end
Return (Detect);

Algorithm 1: Bayesian QCD-based intrusion detection.

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. Algorithm 1 and Algorithm 2 explain how decision making is done using the Bayesian QCD and the Minimax QCD, respectively.

Illustrated in Figure 5 is the variation of detection scores when Minimax 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 151, 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 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 5 samples.

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

//If it's the initial iteration set score to be zero;
if *isempty(*$score$*)* **then**
  $score = 0$;
else
end
//Calculate likelihood ratio;
$L = \log(f_1(x_n)/f_0(x_n))$;
$score \leftarrow score + L$;
if $score < 0$ **then**
  $score = 0$;
else
end
if $score > Threshold$ **then**
  Detect = True;
else
  Detect = False;
end
Return (Detect);

Algorithm 2: Minimax QCD-based intrusion detection.

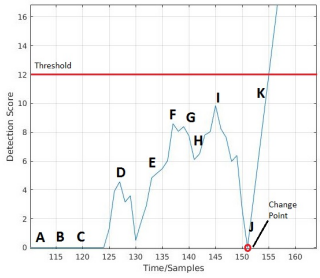

Figure 5. Variation of Minimax QCD scores for the video shown in Figure 1.

4. Experimental Results

We evaluated the performance of the proposed QCD methods using three publicly available unconstrained AA

datasets - Touchalytics [8], MOBIO [13], and UMDAA-01 [7]. 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.) [21]: 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.) [4]: 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 [4] for the exact definitions of f_{map} and f_{dec}.

The PFD-ADD curves, introduced in Section 2, 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 methods proposed by Sui et al. [21] and Crouse et al. [4] 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.

(a) (b)

Figure 6. Sample detected face images from (a) the MOBIO dataset and (b) the UMDAA-01 dataset.

4.1. Protocol

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

4.2. UMDAA-01 Dataset

The UMDAA-01 dataset [7] 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 6(b). As suggested in [7], 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 [5] 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 7 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.

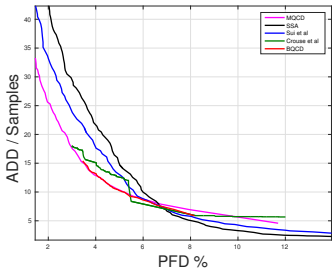

Figure 7. Performance curves obtained on the UMDAA-01 face dataset.

Results on the Touch Data: From each swipe data, a 27-dimensional feature vector is extracted using the method described in [8]. A single class SVM with RBF kernel was used to generate matching scores. Figure 8 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 8, 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. [4] 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.

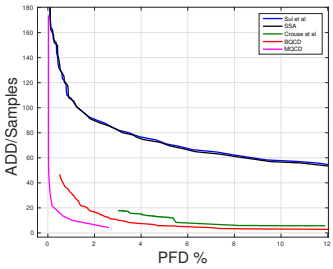

Figure 8. Performance curves obtained on the UMDAA-01 touch dataset.

4.3. MOBIO Dataset

The MOBIO dataset [13] contains videos of 152 subjects taken across two phases where each phase consists of six sessions each. Videos in this dataset are acquired using a standard 2008 Macbook laptop computer and a NOKIA N93i mobile phone (See Figure 6(a)). Following the protocol defined in [16], 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 9 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.

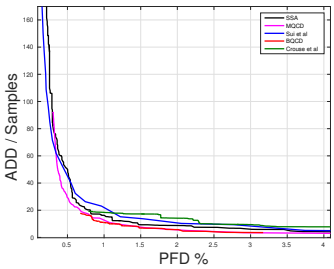

Figure 9. Performance curves obtained on the MOBIO face dataset.

4.4. 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 [8] and a single class SVM with RBF kernel was used to generate match and non-match scores. Figure 10 shows the performance of different methods on this dataset. As before, making a decision based on a single 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. [4].

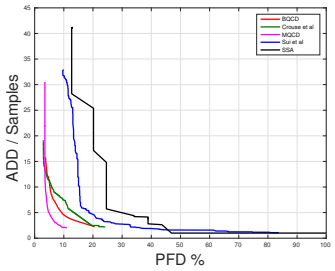

Figure 10. Performance curves obtained on Touchalytics dataset.

4.5. 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. The small margin of error suggests that Geometric distribution assumption is valid for practical applications. However, 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 7, 8, 9, and 10 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 1 is the distribution of detection delay (ADD) for the tests conducted. According to Table 1, 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 1. Percentage breakdown of delay times (in samples) for a fixed PFD of 5% for MQCD. *3% of PFD was used instead.

5. Conclusion

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. It was shown that the proposed method is effective even when there is a considerable overlap between pre and post change distributions.

In the future, we will study the problem of performing QCD in mobile AA systems with constraints on the number of observations and energy.

References

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

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

[3] 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*, pages 199–203, 2006.

[4] 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.

[5] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 886–893, 2005.

[6] C. Deitrick. Smartphone thefts drop as kill switch usage grows, 2015.

[7] 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.

[8] 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*, 8(1):136–148, Jan 2013.

[9] 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*, pages 101–108, Sept 2007.

[10] F. Juefei-Xu, C. Bhagavatula, A. Jaech, U. Prasad, and M. Savvides. Gait-id on the move: Pace independent human identification using cell phone accelerometer dynamics. In *IEEE International Conference on Biometrics: Theory, Applications and Systems*, pages 8–15, Sept 2012.

[11] lookout.com. Phone theft in america, 2015.

[12] G. Lorden. Procedures for reacting to a change in distribution. *The Annals of Mathematical Statistics*, 42(6):1897–1908, 12 1971.

[13] 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*, pages 635–640, July 2012.

[14] E. S. Page. Continuous inspection schemes. *Biometrika*, 41(1/2):100–115, 1954.

[15] 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*, 33(4):49–61, July 2016.

[16] 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.

[17] 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*, pages 1–8, Sept 2013.

[18] A. N. Shiryaev. On optimum methods in quickest detection problems. *Theory of Probability & Its Applications*, 8(1):22–46, 1963.

[19] Z. Sitova, J. Sedenka, Q. Yang, G. Peng, G. Zhou, P. Gasti, and K. S. Balagani. HMOG: A new biometric modality for continuous authentication of smartphone users. *CoRR*, abs/1501.01199, 2015.

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

[21] 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)*, pages 1291 – 1296, 2012.

[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*, pages 207–214, Jan. 2015.

[23] Y. Zhong and Y. Deng. Sensor orientation invariant mobile gait biometrics. In *IEEE International Joint Conference on Biometrics*, pages 1–8, Sept 2014.