Chapter 6
Ensuring Online Exam Integrity Through Continuous Biometric Authentication

Issa Traoré, Youssef Nakkabi, Sherif Saad, Bassam Sayed, Julibio D. Ardigo, and Paulo Magella de Faria Quinan

6.1 Introduction

The last decade has witnessed a growing interest in the area of continuous authentication, with several publications being produced by the research community and diverse products being released by the industry. Continuous authentication consists of verifying repeatedly the identity of a user throughout computing or online session, with the purpose of preventing identity fraud (Traore and Ahmed 2012).

Identity fraud can broadly be categorized in three classes: identity theft, identity sharing, and identity denial. Identity theft occurs when the identity of an unsuspected user is hijacked by a fraudster and used to conduct malicious activity pretending to be the legitimate user. Vehicles for conducting such attacks include phishing, social engineering, and password cracking.

Denial of identity occurs when an authorized individual conducts illegal actions and repudiates such actions when caught. Typically, this would consist of a malicious insider who repudiates malicious actions associated with their identity.

Identity sharing, also referred to as identity gift, occurs when an authorized individual willingly share their credentials with other users, in violation of established policies and regulations. Illegal password sharing can happen, for instance, in the financial industry to circumvent two-man rules, or for paid subscription services such as Netflix.

I. Traoré (✉)
Department of Electrical and Computer Engineering, University of Victoria, Victoria, BC, Canada
e-mail: itraore@ece.uvic.ca

Y. Nakkabi • S. Saad • B. Sayed • P.M. de Faria Quinan
Plurilock Security Solutions, Inc., Victoria, BC, Canada

J.D. Ardigo
Santa Catarina State University, Florianopolis, SC, Brazil

I. Traoré et al. (eds.), *Information Security Practices*,
DOI 10.1007/978-3-319-48947-6_6

A prominent area where illegal credential sharing occurs is online education. With current learning management systems (LMS), students can easily cheat in tests by giving their passwords to others who can take the tests on their behalf. While some Exam Management Systems (EMS) support strong authentication technologies using biometrics, such authentication occurs only statically at login time, this still opens up the door to the possibility for impersonation to occur after the initial login phase.

We propose to address this threat using continuous authentication using a multimodal biometric framework. The proposed multimodal framework combines three complementary biometric technologies: face, mouse dynamics, and keystroke dynamics. All three modalities are collected and processed transparently during the exam without requiring any predefined actions from the test taker.

The proposed framework has been implemented as one the core modules of a new comprehensive exam monitoring platform called ExamShield that has been released recently by Plurilock Security Solutions Inc.

The rest of the chapter is structured as follows. In Sect. 6.2, we discuss and summarize related work. In Sect. 6.3, we present the general architecture of the multimodal biometric framework and its integration in the ExamShield platform. In Sect. 6.4, we discuss the challenges involved in developing our continuous face biometric authentication scheme and give an overview of the approach taken to overcome these challenges. In Sect. 6.5, we make some concluding remarks.

6.2 Related Works

The protection of the integrity of online exams through continuous using biometric technologies is an emerging area of research with relatively few papers (Ahmed and Traore 2011). Furthermore most of the publications, actually, use static biometric authentication.

An example of such line of work has been authored by Ramu and Arivoli by proposing a two-layered approach to address the problem of online exam takers' authentication (Ramu and Arivoli 2013). The two-layered approach combines keystroke biometric authentication and knowledge-based authentication. Although a biometric technology is used, exam participants are authenticated only statically at login time. As mentioned before, this is not enough to prevent cheating from occurring during the course of the exam.

A departure from the above line of work is the approach proposed by Flior and Kowalski who introduced a proof-of-concept implementation of an online exam security system based on continuous keystroke biometric authentication (Flior and Kowalski 2010). In the proposed system, enrolment requires 500 characters collected in a restricted setting (e.g., no backspace or delete is allowed). Furthermore enrolment is performed using fixed text (i.e., predefined text). Similarly, during the exam, verification occurs when 50 keystrokes with no deletion or significant pauses are generated. While the relatively small amount of samples required for enrolment and verification can be considered as a benefit of the system, the restricted nature of these processes will be a significant limitation in real-world deployment. It is not

very realistic to expect an exam to be performed without typos occurring on a regular basis. Furthermore no evaluation of the proposed work has been provided.

Monaco et al. investigated the use of keystroke dynamics and stylometry for continuous authentication of students during online exams (Monaco et al. 2013). Stylometric analysis consists of determining the authorship of a piece of text or document based on the writing style. Like keystroke dynamics, stylometry can be captured transparently using standard keyboard devices. In the proposed work, different studies were conducted using keystroke dynamic and stylometry separately and then by combining both modalities. The combination of both modalities happens at the feature extraction level by concatenating the separate feature vectors into a combined keystroke-stylometry feature vector, which is then submitted to a common classification system. An advantage of this approach over the abovementioned approaches (from the literature) is the use of free text detection, which is crucial to effectively carry continuous authentication. An important limitation, however, is the reliance on a closed-world assumption for authentication. The system relies on a closed population of students serving as basis to train all authorized users. Such assumption is flawed as students cheating in online exams do not necessarily do so with the involvement of other fellow students known to the system. Online cheating may involve sharing credentials with outside individuals totally unknown to the local authentication system.

Fayyoumi and Zarrad developed a prototype for an authentication engine for online exam using continuous face biometric recognition (Fayyoumi and Zarrad 2014). The proposed approach was evaluated by obtaining experts' feedback. Specifically feedbacks were obtained from eight e-learning instructors and 32 students, through a survey using a five-point Likert scale. The proposed examination system includes a question bank which assists instructors in generating randomly different tests for the test takers. Enrolment is performed by capturing and storing images of the user. During the exam, the system tracks the face movement and compares them to the original samples captured during enrolment. A warning is generated in case of suspicion of cheating. A key limitation with the proposed approach is how cheating is characterized. The system relies on facial movement to decide wherever there is cheating or not, which potentially can be a source of large number of false alarms. Furthermore, no evaluation of the performance of the biometric system was conducted. The evaluation was limited as mentioned to the perception of the survey participants mostly on qualitative aspects of the system.

Our proposed framework combines keystroke, mouse, and face biometrics for continuous authentication and does not rely on a closed-world assumption for identity verification. This is made possible by relying only on positive training during enrolment for each of the modalities.

6.3 Online Exam Security: The ExamShield Platform

In this section, we present the ExamShield platform and introduce the general architecture of the underlying multimodal biometric framework.

6.3.1 The ExamShield Platform

ExamShield is a virtual exam center that integrates seamlessly multiple heterogeneous services (multi-biometric authentication, video streaming and recording, exam creation, storage, delivery, and marking). The exam center has been developed as a web portal that can be deployed on the cloud or on premise at the academic institution.

The high-level architecture of ExamShield, depicted in Fig. 6.1, includes the following major services:

1. Exam Management Service provides essential exam management features such as question randomization for different test takers, management of navigation between exam sections, and exam policy enforcement (e.g., exam duration, number of attempt, break time, etc.).
2. Exam Environment Monitoring Service conducts video/audio monitoring of the test taker and surrounding environment using ordinary camera and microphone.
3. Continuous Authentication Service continuously validates the identity of the test taker throughout the exam using a multimodal biometric platform. The platform provides for the first time in an integrated way the following three complementary biometric modalities: mouse dynamics, keystroke dynamics, and facial scans.

Additionally, there is an administrative module which supports technical system administration tasks (e.g., account setup) as well as institutional exam management tasks (e.g., exam scheduling, creating and managing class list and instructors, etc.).

Initially, students are registered to the system by their institutions. Students access their accounts and enroll biometrically once, prior to taking any exam.

Instructors use the system to create and schedule exams. During the exams, instructors access the proctoring panel, where video feeds of the exam participants are displayed. Students are continuously authenticated in the background, and alarms are generated and notified to the instructor through the proctoring panel in real time. Follow-up actions can be taken accordingly by the instructor.

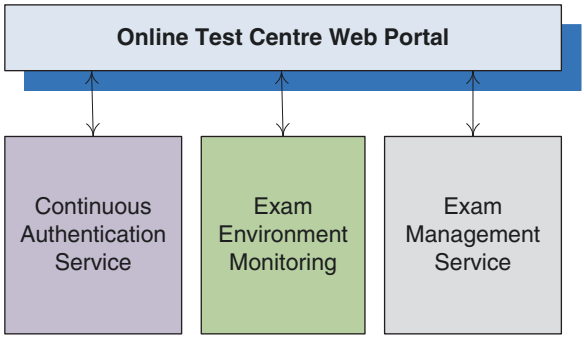

Fig. 6.1 ExamShield high-level architecture

6.3.2 Multimodal Biometric Framework

ExamShield relies on a multimodal biometric framework, through which test takers are continuously authenticated the throughout the exam session from the beginning to the end. The framework involves a combination of several biometric technologies which can be collected and processed transparently in the background without active participation or cooperation of the user.

The multimodal framework integrates three different biometric technologies: mouse dynamic biometric, keystroke dynamic biometric, and face biometric. Mouse and keystroke biometrics are already to be appropriate for continuous as because samples can be collected passively using standard computing devices (e.g., mouse and keyboard) throughout a session without any knowledge of the user. The proposed scheme uses and implements free text analysis and free mouse action analysis models whose theoretical and experimental underpinnings are described in details in Ahmed and Traore (2007, 2014). Interested readers are referred to these publications for details.

Face biometric scans can be collected using standard video cameras, which are currently being shipped with a growing number of computing systems. Facial scans are necessary complement for mouse and keystroke dynamics in order to cover the different monitoring scenarios underlying online exam process. More specifically, while mouse and keyboard may play an active role in written exams, they may be of limited use in exam situations where limited keyboard or mouse interactions are involved, in which cases, facial data could be used to authenticate the user.

However, face biometric has been studied extensively for static biometric authentication; its uses in continuous authentication raise some new challenges since the authentication system has limited control over what the user is doing, which means that there is limited control over the types of samples the application will receive. Hence, it is difficult to capture and analyze effectively biometric samples unobtrusively in a noncooperative environment. Therefore with face biometrics, the recognition must be performed accurately even if images are shifted, or involve different lighting or background, or if the person tilts their head slightly left or right or up or down or angled. This kind of variance between the conditions at enrollment and those at verification times impacts accuracy. Hence, new algorithms must be developed to address the above challenges and ensure effective biometric recognition during online exams. We revisit these issues later in the next section and give an overview of our proposed approach.

6.4 Continuous Face Biometric Authentication

6.4.1 Approach Overview

We designed and developed our continuous face biometric authentication algorithm using local binary pattern and chi-square distance. The model uses only positive training to learn the user's facial features and store the extracted patterns in XML files.

We designed a set of heuristics to improve the accuracy of the system and minimize the false rejection rate.

A major technical challenge in implementing our continuous real-time face recognition over the web was related to the capture and sending of the webcam frames from the user browser to the face recognition server. Existing approaches consist of using communication schemes such as WebSocket to send the captured frames to the face recognition server. However, these do not work for continuous face recognition in a production environment. This is because capturing and sending frames over WebSocket in a continuous setting consume browser resources (e.g., CPU, memory) and result in terminating the WebSocket connection (as in Chrome), or slow down the connection and lose the advantages of real-time authentication (as in Firefox), or even crash the browser and require the user to restart his browser.

The above problem was reported by different developers who were trying to record video stream or send large images or files using WebSocket. Most of the suggested solutions focus on decreasing the video frame rate and connection time. However, this is not possible in our application because the recorded frames are used both for face recognition and user authentication. In addition, an online exam can take up to 4 h and in some cases more than that. To solve this problem and get beyond the current limitations, we took the following steps:

- Use WebP image encoding and avoid using PNG and JPEG image encoding (some browsers do not support WebP).
- Adjust the frame rate and image resolution based on the browser support for a particular image encoding.
- Send binary image not base64.
- Implement a fault-tolerant technique to detect WebSocket connection drop by the browser.

While the matching performances of the above scheme are excellent, its success depends on facial feature being tracked effectively. The OpenCV library is the reference framework for computer vision and face implementations. However, tracking a face in video stream is not an out-of-the-box feature in the OpenCV library. This is because there is no single face-tracking algorithm that can serve different applications. For example, a face-tracking algorithm in video game console is not appropriate for other applications. To address the specific challenges of continuous face biometric authentication, we implemented initially two new techniques to support face tracking. The first technique consists of a motion detection algorithm that calculates the difference between two consecutive video frames and, based on a predefined threshold, decides if a motion exists in the video stream or not. The second technique relies on using existing OpenCV face detection algorithms, and then after the initial detection, it performs a template matching to detect the face template in the new video frames.

These two techniques yielded acceptable performance in laboratory and offline testing environments. However, they did not yield the same performance in online testing with live subjects performing real-world tasks. This is mainly because of the restrictions we have on the video stream. These include the video resolution,

which can be extremely low due to the diversity of exam participants (e.g., low-end Internet connectivity in some countries, heterogeneous platforms); the frame rate per second; and the fact that the system is running in uncontrolled environment, which is a typical characteristic of continuous authentication.

Likewise the existing tracking algorithms publicized through OpenCV do not scale/perform well in real-world environment confronted with the need for the flexibility inherent to continuous authentication. Through thorough search, we could not find in the literature any tracking algorithm that addresses our exam environment constraints (e.g., 5 fps and 320×240 resolution, webcam, and uncontrolled lighting). To address these limitations, we had to make changes to the way the recognition algorithm works and to work with available face frame and do not require a specific number of frame to take the decision. With our new algorithm, even a single frame with one face can be used for recognition, while previously at least 300 frames were required.

6.4.2 Evaluation and Observation

To evaluate the performance of our continuous face recognition system, we divided the evaluation process into three main phases. The first and the second phases focused on evaluating the recognition accuracy, while the third phase focused on evaluating the system in the production environment. In the first phase, we evaluated the detection accuracy of our face recognition algorithm with respect to positive training and novelty detection. To evaluate the detection accuracy when using only positive training, we used existing benchmark facial recognition datasets that are commonly used for evaluating static face recognition algorithms. We used the following three datasets: the AT&T Face Database, the Yale Face Database, and the extended Yale Face Database B. Our face recognition system yielded an accuracy between 91.32 and 94.71 %. These results are very encouraging considering that the algorithm uses only positive training. Most existing face biometric depends on both negative and positive training.

In the second phase of our evaluation, we recorded video streams from 11 subjects. Each subject has to visit our face recognition web application. The video streams were captured using WebRTC and transmitted to our server using web sockets. The server is implemented in python; we used Twisted and Autobahn as our network framework. All the video and image processing are handled by the OpenCV library. About five or six video streams for each subject were recorded. Each video stream is 10–15 min length. These video streams were recorded using a webcam. We used the first 3 min of video data for training. So only 3 min of the 50 min of each subject was used for building the subject face signature. Finally, we merged all the recorded video streams into one big video file and used this file to evaluate our continuous face recognition algorithm. The accuracy of the system in this experiment was 100 %. The system was able to always distinguish the legitimate subject from the imposter subject. While in static authentication our best result was 94.71 %,

in continuous authentication our result was 100 %. Such difference in accuracy is mainly due to the fact that our algorithm was designed for continuous authentication. So it was able to take advantages of the huge amount of data (300 face samples per minute) it has for training and verification in comparison to the limited number of face samples (e.g., 20 face samples) used in static authentication.

The last phase of our evaluation focused on evaluating the system in the production environment. The face recognition server was deployed on the cloud. One instance was deployed on amazon cloud on the west coast, and another instance was deployed on a private cloud hosted by Plurilock Security Solution Inc., in Victoria, BC, Canada. In collaboration with different institutions, students from Canada and Brazil connected to the ExamShield server to perform live exams over several sessions. The students were invited instructed to create their facial signatures prior to taking online exams. The face recognition system was able to record the exam sessions for all the students, perform face recognition and verification in real time, and generate alarms in real time to notify exam proctors. Alarms were generated when a student was taking an exam on behalf of another student, when the student leaves his chair during the exam, or when several students were working on the same exam together. During the production evaluation, most of the reported problems were related to technical problems such as memory leak, connection drop, etc. All these pure technical issues were handled and fixed. The most interesting issues that were reported during the production testing and affected the face recognition functionalities were related to the environmental/external conditions that appear in the exam session. For instance, a major change in the lighting conditions, such as turning the light off during the exam or changing the location of the desk lamp, can badly affect the recognition accuracy. These observations show the need for a real-time adaptation technique to mitigate the effects of the extreme external factors in the exam environment. This will be one of the main focuses of our future work.

6.5 Conclusion

Continuous authentication is an emerging technology which is proving to be appropriate in handling a variety of security threats. Concrete applications range from forensic analysis, detection of insider threat and session hijacking, and various forms of illegal identity sharing. Cheating in online exams falls in the latter category.

This chapter introduces a multimodal biometric framework combining for continuous authentication of online test takers. The framework represents a core module of the ExamShield platform, which is a new online exam monitoring system. In addition to continuous authentication, the ExamShield platform provides live video streaming and recording of exam environments and essential exam management services. The different biometric modalities have been evaluated separately using offline datasets. The biometric framework is currently being used in production in the ExamShield platform yielding very encouraging results.

It is important to highlight that while the biometric framework involves multiple biometric technologies which are complementary, each of the modality is processed separately, and the outputs are presented through separate authentication events displayed using a common dashboard. Likewise, the framework may not technically be considered as full multimodal scheme, as there is no fusion of the outcome of the separate modalities.

In our future work, we plan to address such gap by developing a fusion scheme that will combine the three biometric modalities involved in the framework (i.e., mouse, keystroke, and face) and generate a combined and unique score for overall decision-making.

The effectiveness of a multimodal scheme depends on the appropriateness of the underlying fusion technique used to combine the outcome of the separate modalities. Traditional fusion techniques rely on the availability at the time of the fusion of the separate information being fused. More specifically, the outputs of the separate biometric modalities must be synchronized.

However, synchronizing such a process is not appropriate for continuous authentication as this will delay some of the modalities, which leads to longer verification time.

Our goal is to develop an asynchronous fusion model based on the sequential sampling theory that will allow making a trade-off between accuracy and authentication delay, which is needed in continuous authentication (Ahmed and Traore 2011).

References

Ahmed AAE, Traore I (2007) A new biometrics technology based on mouse dynamics. IEEE Trans Dependable Secur Comput 4(3):165–179

Ahmed A, Traore I (2011) Dynamic sample size detection in continuous authentication using sequential sampling. In: Proceedings of annual computer security applications conference (ACSAC), 5–9 December 2011, Orlando, FL, USA

Ahmed AAE, Traore I (2014) Free text recognition of keystrokes. IEEE Trans Cybern 44(4):458–472

Fayyoumi A, Zarrad A (2014) Novel solution based on face recognition to address identity theft and cheating in online examination systems. Adv Internet Things 4(3):5–12. http://www.scirp.org/journal/ait, http://dx.doi.org/10.4236/ait.2014.42002

Flior E, Kowalski K (2010) Continuous biometric user authentication in online examinations. In: 2010 Seventh international conference on information technology, pp 488–92

Monaco JV, Stewart JC, Cha SH, Tappert CC (2013) Behavioral biometric verification of student identity in online course assessment and authentication of authors in literary works. In: IEEE 6th international conference on biometrics, BTAS

Ramu T, Arivoli T (2013) A framework of secure biometric based online exam authentication: an alternative to traditional exam. Int J Sci Eng Res 4(11):52–60

Traore I, Ahmed AAE (eds) (2012) Continuous authentication based on biometrics: data, models, and metrics. IGI Global. ISBN: 978-1-61350-129