UNIVERSITY OF MOLISE

Department of Biosciences and Territory

Master Thesis in Software Systems Security

EDISON: a Time-Oriented Context and Application Adaptive Continuous-Authentication Framework

Author:
Dr. Jonathan Simeone

Supervisor:
Prof. Stefano Ricciardi

Biometric Systems

December 15, 2022

Do things that build confidence

Abstract

Continuous Authentication systems possibly represent the future of *user authentication* and research is continually moving forward through the use of new biometrics and innovative techniques for capturing, fusing and evaluating them to verify user authenticity. However, there are some perspectives that seem to have not yet been fully explored, such as the potential of *context awareness*, and *biometrics reliability* assessment, as well as the exploitation of *temporal aspects characterizing the authentication flow* for the decision task, which is commonly based on instantaneous evaluations instead. To this regard, the goal of the current work is to investigate the potential of these aspects to include them in *Edison*, a *Continuous Authentication Framework*, a novel unified approach capable to exploit these new elements by abstracting itself from specific types of biometrics or biometric systems architectures. More precisely, this research proposes a continuous-authentication approach that is *fully generalizable*, *biometric-independent*, *context and application adaptive* and that exploits *time series analysis* in order to associate an *authentication level* to the user (which may be sufficient in a particular context and not sufficient in another) instead of simply providing a binary authentication decision. In order to concretely evaluate the validity of such an approach, an experimental framework has been implemented and subjected to a preliminary *testing*, which confirmed its potential as well as showing some aspects deserving further development in the future.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Application Context

Authentication is a way of ensuring that an individual is authorized to gain access to a specific location or allowed to perform specific operations. We can look at what authentication systems offer in companies to protect business assets, in computer systems to protect sensitive information, but also in our homes and even in our cars to ensure the authenticity of an individual in order to prevent theft or general damage of any kind. Since this need of security has been felt, there have been several evolutions of the methods of authenticating a user. We started from having a simple password to moving to credentials of increasing complexity, and then to considering the idea of using not something to remember to authenticate, but something that belongs to us, such as our face, our voice, our eyes, and so on. The evolution has not stopped there: ingenious new biometrics have been discovered, others has been merged, and research is continually pushing new approaches in order to improve already existing techniques. As result of such studies, sophisticated multi-biometric systems are now available that offer an excellent level of security. In particular, thanks to novel techniques for capturing, processing and eventual merging multiple biometrics, main problems affecting simpler single-biometric systems, such as biometric cloning, have been mitigated. Finally, dropping these results into a *Continuous Authentication* operating mode, made the authentication task even more effective. In fact, we have moved beyond the simple notion of one-time authentication to proceed to techniques can be adopted to obtain a Continuous Authentication (from now *CA*) from one or several biometric channels in order to check over time whether the user's genuineness can still be verified. In this way, even if a way to cheat one or several biometrics could be found, the task would have to be repeated during an entire authentication session, which would make an effective intrusion of an impostor very complex.

1.2 Motivation and Objectives

In this research, we aimed at exploring the territory of CA systems in an attempt to take another step toward to improve the effectiveness of such systems. A few recent surveys on this topic, indeed, provided some significant insights.

First, it seems that a key element that could be a breakthrough in the analysis of biometric acquisitions has not yet been sufficiently considered. In particular, the *temporal factor* that correlates different acquisitions over time has not yet been thoroughly investigated: typically, the other works dwell on specific biometrics and new approaches to exploiting them, and it all boils down to a simple score that authenticates or does not authenticate the user for the biometric(s) being analyzed at a given time. So, what seems to be missing is a focus on what can be defined as an *Authentication Flow*, *i.e.,* a set of local decisions (obtainable by any approach for any biometric) that can be subjected to such a temporal analysis to finally obtain an *Authentication Level*. This step is not to be underestimated: having an authentication level moves us away from having a simple binary decision, based on which user will be authenticated or not; rather, an authentication level would allow the user to be authenticated for one specific context, and *"de-authenticated"* in another. In other words, a user can *maintain* an authentication level sufficient for one task or one place, and be de-authenticated for another one if his authentication level is too low.

Here we come to another interesting point, which is precisely the context factor. If we could equip a CA system with a *Context Awareness* mechanism, in fact, we could pursue all that we have just described. If a CA system is *Context-sensitive*, in fact, it will be able to weigh its decisions based on what the user is doing, where he is and what surrounds him.

In addition, it was noted that there was no particular focus on the *reliability* of an authentication channel. In fact, a reliable CA system, should notice *anomalies* coming from a channel in order to block it as long as the anomalies continue to occur. We are talking, then, about equipping a CA system with *Reliability Assessment* mechanisms.

As it turned out, the purpose of this thesis work is not to find new biometrics or to improve existing biometric capture or processing techniques. The intended contribution of the research conducted is to offer a novel way of putting into a new perspective the data that a continuous authentication system collects through his sensors, to offer a new way of evaluating and enriching that information, and to encapsulate all these features in a novel *CA Framework* that will be absolutely *generalizable* and not specific to one particular type of biometric or architecture. The framework, therefore, will represent a kind of guideline that will hopefully help the new generation of CA systems to be more effective thanks to a new way of employing *temporal*, *context-related* and *reliability-related* informations.

Such an approach can be imagined as applied in a corporate setting, where there may be several sensors in some top security rooms, such as cameras and termoscanners, which can be leveraged to achieve continuous and effective authentication of authorized employees. Another example may be the required authentication in a smart car: if a car is equipped with an effective CA, in fact, it could alert the authorities the moment the driver is not recognized as authentic. Finally, the most effective example is to imagine a framework like that actualized in a smartphone that mount numerous sensors and that will surely increase in the future. If we combine the availability of sensors, and consequently the possibility of having multiple biometrics, with a device that now constantly accompanies us throughout our days, we potentially gain the ability to have an effective continuous authentication mechanism available at all times, which can be used to protect the smartphone data itself but also to offer a bridge for external applications that will use the current level of authentication to authorize or not the user in whatever context he is operating. If, for example, a user is at home and the framework running on the device is aware of this, the smart home's security mechanisms can use this information to take some decisions. If the user moves to the office, the security mechanisms of the smart office can keep the user authenticated for some rooms by leveraging a connection to his personal device. If the user rents a car that can only be driven by him, the smart car mechanisms can use the smartphone authentication info complementarily to the car security system informations to keep authenticated the driver.

Based on all the considerations made, and based on what are the aspirations of this thesis work, the following research questions were posed:

RQ1: *Is it possible to build a CA Framework by abstracting its input and thus not being tied to specific biometrics or architectures?*

RQ2: *Can a Temporal-oriented Analysis of an Authentication Flow be more effective in maintaining the authentication of a genuine user than considering a single instantaneous decision?*

RQ3: *Can a channel Reliability Assessment mechanism be useful to maintain an high authentication rate in a CA system?*

RQ4: *Can a Context Awareness mechanism make a CA system less invasive in a normal condition and more secure in high-security settings?*

1.3 Results

Having clear in mind the objectives of the work, the framework was designed in all its components thanks also to the preliminary study of the state of the art that made it clear how a continuous authentication system generally works. The first step, in fact, has been to model a generic CA system in order to updating it with the insights that are the subject of this thesis work. Then, as final result, we modeled a new block diagram that fully represents the framework and its operation. Subsequently, in order to give validity to the work, a preliminary experimentation was conducted to confirm the potential of the approach and to highlight any critical issues. In order to carry out the experimentation, therefore, a possible implementation of the framework was designed and implemented, which, for reasons that will be clarified in the next chapters, took the form of an Android mobile application. The results showed that the approach definitely has the hoped-for potential, but also some weaknesses that can be improved with future developments. In particular, the potential of temporal flow of authentication and the effectiveness of context detection and reliability assessment mechanisms were questioned through the above-mentioned research questions. In addition, at the operational level, we wondered whether it was actually possible to generate a Continuous Authentication Framework that could exploit these concepts by also abstracting from specific biometrics and architectures. The results showed that authentication flow is definitely effective (although with some improvements to be made regarding the speed with which an impostor will then be recognized) compared to a point decision based on a simple score; we also ascertained that an intelligent reliability mechanisms can keep authentication performance high and that context detection mechanisms allow for maximum flexibility but also maximum security. Furthermore, the concrete implementation of a mobile app generated by the designed framework demonstrated that it was indeed possible to define an approach that did not depend on specific biometrics or specific architectures. In fact, in both the choice of architecture and biometrics we did not notice any limitations from the framework as designed.

1.4 Thesis Structure and Organization

The next chapters of this thesis is structured as follow:

- *Chapter 2*: presents the state of the art and related works;
- *Chapter 3*: presents the framework with its components with a detailed description;
- *Chapter 4*: presents an example implementation of the framework born to conduct an experimentation of the novel approach;
- *Chapter 5*: describe the conducted experimentation and its results;
- *Chapter 6*: describe the results of this thesis and provides directions for future work.

Chapter 2

Background and Related Work

There are several ways to approach the CA problem. An overall picture of the approaches to this topic available so far in the literature, can be found in the study by Ayeswarya and Norman "A survey on different continuous authentication systems" [1]. In their work, the authors provide a thorough overview of existing continuous authentication methods, including their strengths and limits. In general, research has explored the usage of many different types of biometrics; an example is the work "Continuous authentication by free-text keystroke based on CNN and RNN" by Lu *et al.* [2], in which the authors described an approach able to authenticating users via their keystrokes when they type free text. In the work by Siddiqui *et al.* [3] "Continuous Authentication Using Mouse Movements, Machine Learning, and Minecraft", mouse motion as videogame controller is used for discriminating between specific users moves. A number of related works involve the use of smartphones to both capture and process biometrics. An example is from Frank *et al.* and their work "Touchalytics: On the Applicability of Touchscreen Input as a Behavioral Biometric for Continuous Authentication" [4] that describes a CA classifier that recognizes users based on the way they interact with the smartphone via its touchscreen. Moreover, of particular relevance, are multimodal biometric authentication systems that succeed in mitigating the problems related to single biometric systems. One example is the work "Continuous user authentication using multi-modal biometrics" by Saevanee *et al.* [5], where the authors proposed a novel text-based multimodal biometric approach utilizing linguistic analysis, keystroke dynamics and behavioral profiling. Another example is the case of the work "KeyGait: Framework for Continuously Biometric Authentication during Usage of a Smartphone" from Trojahn and Ortmeier [6] in which the authors proposed an approach that exploit keystroke dynamics and movements of the smartphone to authenticate users. In general, there is a plethora of works which take advantage of smartphone sensors. Abuhamad *et al.* [7], in their work "Sensor-based Continuous Authentication of Smartphones' Users Using Behavioral Biometrics: A Contemporary Survey", reported more than one hundred and forty recent behavioral biometric-based approaches for continuous user authentication that uses smartphone sensors. There are also CA system implemented on wearables. An example is the study "WACA: Wearable-Assisted Continuous Authentication" by Acar *et al.* [8] in which authors proposed a Wearable-Assisted Continuous Authentication approach that exploits authentication data of keystroke dynamics acquired through the built-in sensors of a smartwatch. Another related work is "Smart watch based body-temperature authentication" from Enamamu *et al.* [9] that proposes BT-Authen, a CA approach that uses to authenticate users galvanic skins response to extract the body temperature information via the wearable sensor. But the wearable devices used for such purposes are not just smartwatches. This is the case the work "Continuous Authentication With Touch Behavioral Biometrics and Voice on Wearable Glasses" by Peng *et al.* [10] in which is described GlassGuard, a CA system running on smart glasses that uses six typesof touch gestures and voice commands to discriminate owners and impostors. Research is continually evolving on possible data that may represent a new biometric; an interesting example is the work "Real-time electrocardiogram streams for continuous authentication" by Camara *et al.* [11] in which is described a continuous authentication approach based on users real-time ECG signals as continuous data stream. Many of the modern CA systems, moreover, use the potential of the IoT world. As evidence of this, Manzano *et al.*, in their work "Leveraging User-related Internet of Things for Continuous Authentication: A Survey" [12], made a review of a large number of research papers about IoT-based CA systems and their main components.

The analysis of the state of the art in CA has been the foundation of the approach and framework proposed in this thesis, providing possible insights on the areas requiring improvement of novel solutions.

Specifically, the first insight came from the paper "Continuous authentication using biometrics: An advanced review" by Dahia *et al.* [13], in which are described various biometrics that can be implemented in a CA system. By looking at their reports, we were able to observe the advantages and disadvantages of such biometrics. From here we wondered if a possible advancement of the state of the art was possible by going to work on new biometrics or new approaches to evaluating them or do better trying to go further and hypothesize an analysis that takes into account not individual results, which indeed can have disadvantages for specific biometric-related problems, but a set of results that, under optimal conditions, might give a clearer idea of what an individual's authentication status is. But what element could define how to correlate a set of individual authentication states? Well the idea for an answer was given in the same article, which mentions: "It is noticeable that many works do not address temporal integration at all, and instead approach continuous authentication as a series of independent traditional authentications. This choice discards valuable information regarding the validity of previous authentications, that is, it is very unlikely that a genuine user verified with high confidence is immediately replaced by an imposter". This statement was of fundamental importance: a key factor that apparently has not yet been given the right weight is the temporal factor that can bind different acquisitions of biometrics. As a result, we could see how a novel and as yet insufficiently explored approach necessarily had to be based on a *temporal analysis* of a range of recent *authentication states*.

Another paper inspiring this the proposed approach was "Behavioral Biometrics for Continuous Authentication in the Internet of Things Era: An Artificial Intelligence Perspective" by Liang *et al.* [14]. In fact, in this paper, an *anomaly detection* system is cited as a possible tool to effectively achieve an authentication system based on behavioral biometrics. The anomaly detection is also key in "An Empirical Evaluation of Online Continuous Authentication and Anomaly Detection Using Mouse Clickstream Data Analysis," by Almaki *et al.* [15]. In this study, an empirical evaluation of continuous authentication (CA) and online anomaly detection (AD) based on the analysis of mouse click stream data is presented. As we can see, an anomaly detection system in the cited literature is identified as a mechanism by which one can actually notice a variation in a signal and use this information to, eventually, unauthenticate a user. So, we wondered whether it might not be possible to exploit an anomaly detection system not as a mechanism to unauthenticate a user, but rather to try to keep the user authenticated when it is noticed, through appropriate anomaly detection mechanisms, that the biometric channel that is attempting to throw out the genuine user is not stable.

Another interesting work is found in Brosso *et al.* [16], with their study "A Continuous Authentication System Based On User Behavior Analysis". In their work, they developed a biometric system that captures information about the environmental *context* to interpret user intention and analyze his behavior. In this case that context information is used complementarily to generate a behavioral biometrics, rather than just using interesting information such as context information only for a biometric definition. We wondered whether it could be worth using it to understand something more important, namely, the *level of criticality* at which his authentication state need to be evaluated. Indeed, based on context information, one can define whether the context of use of the authentication system needs more flexibility or more stringent security requirements. Thus, an authentication system could be less intrusive if it is not needed from the system use scenario and more secure in critical contexts.

Further food for thought was given to us by how some studies usually involve the choice of particular biometrics and their fusion through particular approaches in order to improve effectiveness. One such study is by Azzini *et al.*: “A fuzzy approach to multimodal biometric continuous authentication” [17]. They developed a multimodal biometric system focusing on the possibility of using a fuzzy controller to conduct a multi-modal continuous biometric authentication using face and finger biometrics. In particular, if a sufficient frame rate of the analysis is guaranteed, a face trait can be used to ensure a continuous control of the users’ faces until a threshold is respected. If not, a fingerprint authentication is requested with also a new face acquisition and then a fuzzy controller computes two fuzzy variables, defined by the two biometrics template matching scores, and returns a trust value for the user. From what can be seen, the approach described by this research work, with its modes and configurations, are specific to the particular biometrics used. We wondered whether it might not be possible instead to create an approach that could be generalizable and scalable without having to be limited by specific biometrics.

In addition to aiming at a biometrics-independent approach, we thought about whether this goal could be achieved in an architecture-independent fashion as well. In fact many works, as well as the last one mentioned, involved a fixed capture station, while a lot of other works were in a mobile context using smartphones. An example is from Crouse *et al.* with their work “Continuous authentication of mobile user: Fusion of face image and inertial Measurement Unit data” [18] that describe a face-based continuous authentication system outlining a method for compensating for camera orientation and, consequently, the orientation of the face picture on a mobile device utilizing inertial data. In Abate *et al.* with the work “I-Am: Implicitly Authenticate Me - Person Authentication on Mobile Devices Through Ear Shape and Arm Gesture” [19] te authors describe a multibiometric system based on the discovery that the instinctive action of answering a phone can be utilized to collect two distinct biometrics, specifically the ear and arm gesture, which are complementary due to their, respectively, physical and behavioral character. The last example is from Buriro *et al.*, that in their study “Please hold on: Unobtrusive user authentication using smartphone’s built-in sensors” [20] described a completely unobtrusive user authentication method based on minute motions of the user’s hand(s) following smartphone unlocking.

As we observed, analyzing the state of the art in Continuous Authentication systems, several interesting insights emerged. In particular, there is the possibility of seeing both anomaly detection and context detection mechanisms from different aspects. Also, it seems that the temporal factor of different biometric acquisitions has not been sufficiently investigated. What is more, there seems to be no approach that is sufficiently generalizable that it is not tied to specific biometrics or architectures. These consideration led us to the idea of unifying these key points together through a novel Framework, fully exploiting the temporal flow of biometrics events and featuring a scalable, generalizable, context-sensible and anomaly-aware approach to CA.

Chapter 3

The Edison Framework

3.1 Approach Key Points

Edison (gEneralized aDaptive contInuouS authenticatiOn eNgine) represents a novel framework for Continuous Authentication. The aim, in particular, is to offer an intelligent, scalable and effective approach to the problem of continuous multi-biometric authentication. Several approaches to such problems, as we have seen, can be found in the literature. Yet, as we have seen, these approaches typically involve the use of different biometrics and the creation of methods and algorithms to optimize their accuracy.

Edison's aim, on the other hand, is to abstract itself from specific hardware and biometrics in order to offer an approach that can be modeled on any hardware and biometrics, offering a scalable, reliable, context-adaptive and effective system designed for high usability and security. With a view to the future, when offices, homes and even our smartphones will be equipped with an increasing number of sensors, a scalable and generalized framework, capable of adapting to variable operative and applicative conditions, could be a considerable advantage. The work behind Edison is inspired by some considerations as to what is expected from a continuous authentication framework that fulfills the idea described so far. In particular, following an analysis of what is the state of the art reported in the previous chapters, some considerations were deduced:

- A multi-biometric CA framework does not have to provide a set of usable biometrics, nor necessarily specific configurations for each of them. The concept of biometrics must be abstracted in order to capture the key elements common to any biometric channel;

- An evolved CA framework must be equipped with context detection techniques that allows the system to adapt the authentication mechanism to the current scenario while increasing both usability and the level of security;
- The point-in-time verification of a biometric score on a threshold may not be effective when the biometric signal is affected by some momentary disturbance; on the contrary, a temporal analysis over several authentication states may offer better performance;
- A biometric channel may present anomalies. A state-of-the-art system must realize this and protect the authentication flow from anomalous channels until the signal is stabilized;
- A CA system has the possibility of acquiring many probes over time. Considering that some of these may be qualitatively excellent and at the same time also return high authentication scores, it may be useful to provide a continuous enrolment mechanism.

Based on these considerations, the key points of the framework were identified. In particular:

- *Biometric Abstraction in an authentication state*: any biometric, be it physical or behavioral, must be abstracted by extrapolating what are the important elements for an authentication, in particular: *the probe*, *i.e.,* the acquired data (regardless of its nature), *signal quality*, *i.e.,* a value representing the goodness of the signal and *the score*, *i.e.,* the authentication score that the algorithm (computational, statistical or machine learning related) gave to the acquired probe comparing it with the user template;
- *Context detection mechanisms*: the detection of the context makes it possible to offer maximum usability of the system by avoiding explicit (and therefore invasive) authentication requests when the scenario of use is not critical. Conversely, if the scenario is critical, the system increases its security level by possibly requesting explicit authentication;
- *Decision engine with temporal analysis on authentication flow*: not only separate authentication events need to be evaluated, but also a temporal authentication

flow has to be analyzed in order to make a final decision. Authentication is enriched with new states over time, all of which have a certain importance depending on how recent they are and any other eventual weights defined in the framework configuration;

- *Dynamic reliability assessment on channels*: an anomaly detection mechanism can contribute to the reliability of a CA framework. Only states that do not come from a channel that is generating anomalous states will be included in the authentication flow. It is not simply a quality control of the signal (which is still present by preventing the probe acquisition or weighting the score), but a true analysis of the states generated over time to check for anomalous fluctuations in the score series of a specific channel. It must be considered, however, that an anomalous channel must be always recoverable when the anomaly in the signal ceases;

- *Continuous enrollment*: a problem that plagues many biometrics is the missing of persistence of the biometric traits over time. A continuous enrolment system can follow the changes of a subject by frequently and automatically updating enrolment templates when the probes obtained during authentication have a high quality and high score.

3.2 Evolving a generic Continuous Authentication System

Having clarified the key points of the approach, a modeling of the framework was carried out. In a totally abstract manner, a generic ideal architecture of a multimodal continuous authentication system was modeled. In particular, it provides for the input various biometrics that will be processed by an engine that applies a precise fusion strategy dependent on the biometrics in use. The result of this fusion generates a score that will be compared with a threshold. Following this comparison, the user is either authenticated or rejected. Fig. 3.1 shows a model as described.

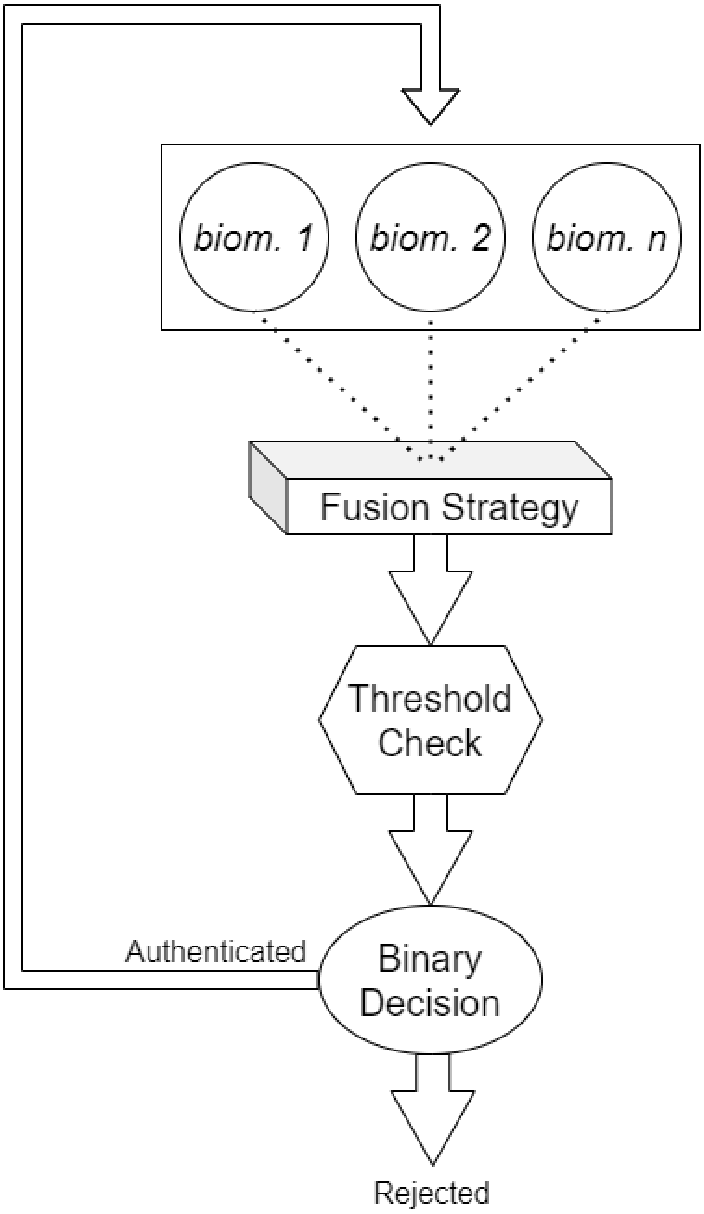

FIGURE 3.1: A generic multimodal CA system.

Obtained a base model, the key concepts described earlier for the framework were added (Fig. 3.2). In particular:

- abstraction from biometrics has been introduced;
- authentication informations from a each biometric has been enriched with context and channel anomaly information;
- informations from the biometrics enriched with context and anomaly detection elements are set to be processed by a decision engine that first uses a local decision controller to compute a local decision score and next performs a temporal

analysis on an authentication flow consisting of a certain number of recent received states. The context informations allow set the more suitable threshold and, in addition, if the local decision has an optimal value the template is updated;

- finally, a level of authentication is calculated to be compared with the dynamic threshold that varied according to the criticality of the context in which the user is located.

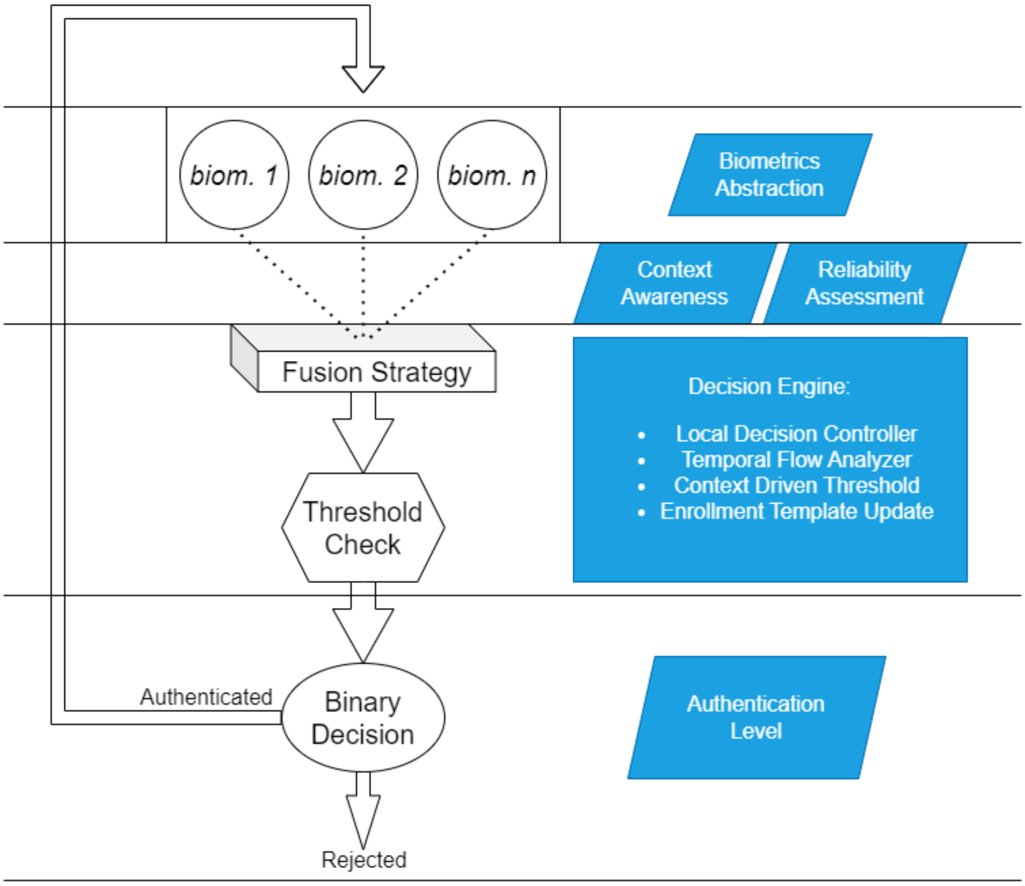

FIGURE 3.2: An Edison-oriented multimodal CA system.

3.3 Framework Design

Having in mind what a final Edison-oriented system might look like we started to design it going to imagine what the flow of data within such a system would be. Is now described the high level data flow that start from a single biometric channel all the way to a level of authentication. In particular, let assume that we have a generic biometric channel *BIOn* on which an enrolment has been carried out and that we are in a continuous authentication session:

- Quality and contextual information on a possible acquisition of the *BIOn* channel is gathered;
- If acquisition is possible, the biometric probe is acquired, a score is calculated using a specific approach and the signal quality information is also stored;
- A local decision is calculated by a controller that can also weights the score result using the quality information; if the score is optimal and the quality equally excellent, the enrolment template is updated;
- The local authentication state triggers anomaly checks, performed by analyzing the current and old authentication values on the specific *BIOn* channel. If the channel is anomalous, it will be ignored until it is stabilized;
- If the channel is not abnormal, the current authentication state is added to an authentication flow that includes any recent *BIOn* state as well as other biometrics. An authentication level for the user is calculated from the authentications flow from the temporal flow analyzer;
- Based on the context information, the threshold is set at a certain level. Depending on the result of the time analysis, there will be a certain level of authentication, which may be, for example, sufficient for a low-security context and insufficient for a critical context;
- Depending on the result of the comparison with the dynamic threshold, the system will either maintain the authentication state by proceeding with new *BIOn* acquisitions or require explicit authentication over a secure channel. In fact, unlike common authentication systems that usually return authenticated/unauthenticated decisions, Edison wants to describe the concept of maintaining an authentication state and reach, unfortunately, a *"de-authentication"* state that represents, precisely, a transition from an *authentication state* to a *non-authentication* state, i.e. we talk about a loss of an authentication state.

Having defined the key concepts of the framework by evolving a generic CA system and having described its operation in the main phases, it now possible to generate a complete model of the framework as designed that is shown in Fig. 3.3. In particular, Edison can be represented by a block model in which there are the described components that give rise to our continuous authentication framework. In particular, assuming we want to implement a CA system using Edison approach, we must follow the following framework steps:

- *Choice of biometrics* (physical and/or behavioral) and implementation of enrolment, verification and quality assessment mechanisms;
- Abstraction of each biometric in a *Authentication State* that will contain useful information such as a signal quality value (that can be used to weigh the score) and the score obtained from a specific approach or even from the fusion of several biometrics;
- Implementation of a *Context Detection* mechanism that can provide useful information by analyzing environmental parameters and user activities and that can allow the system to adapt to different levels of criticality;
- Implementation of a *Reliability assessment* mechanism that performs a dynamic analysis on the data coming from the different channels in order to block them until they return stable values over time, with a controller that also checks at a given time whether at least one stable channel is available;
- Implementation of a *Decision Engine* with *Configurable Security Parameters* capable of:
    - Using a *Local Decision Controller* that can generate *Local Decisions* using all the necessary info contained in the incoming *Authentication State*;
    - *Updating the template* in case of optimal probes acquired both in terms of score and quality;
    - Maintaining an *Authentication Flow*, *i.e.*, a temporal queue of Authentication States;

- Analyzing the Authentication Flow using a *Temporal Analysis* in order to generate an *Authentication Level*;
- Using the Authentication Level and the *Level of Criticality* detected by the context detection for maintaining the authentication or de-authenticating the user requesting *Explicit Authentication* through a secure channel.

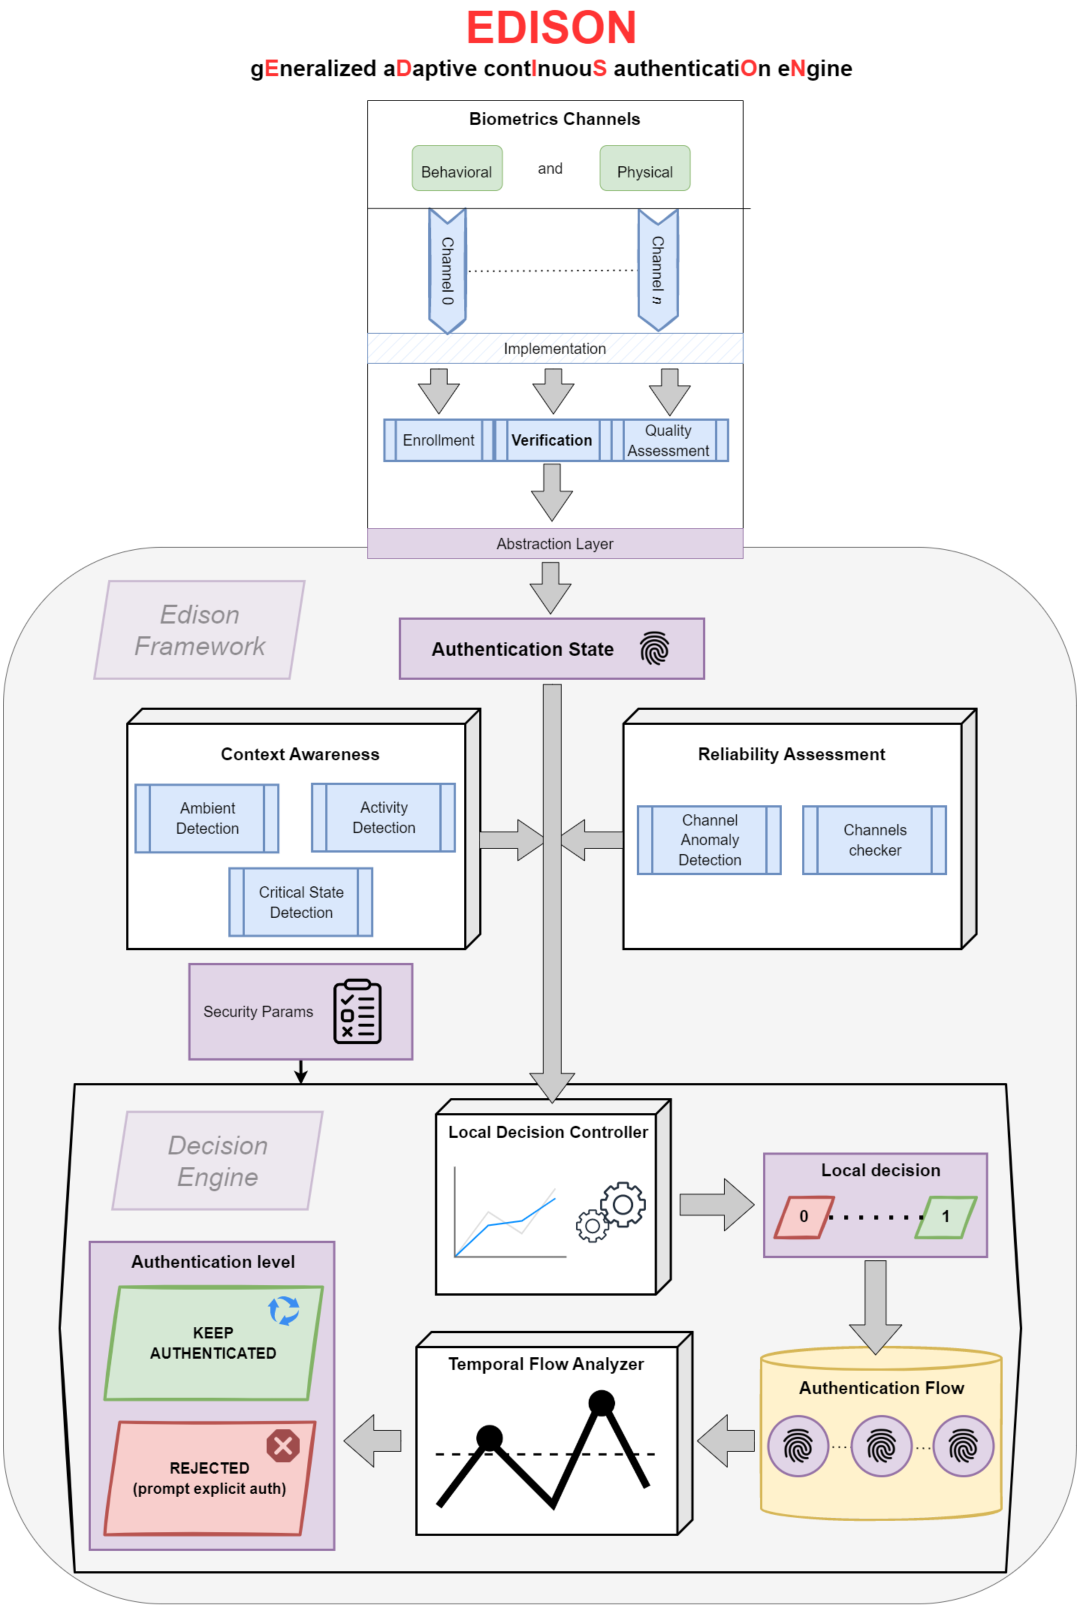

FIGURE 3.3: Diagram of the Edison framework approach.

The individual components of the proposed approach are now described in detail.

3.3.1 Biometrics Channels

It represents the input of the authentication system and allows the implementation of the chosen biometrics that can be either physical or behavioral. At this stage, it's even possible to merge different biometrics with any approach. Furthermore, it is also possible to use channels of the same type, i.e. to have acquisitions of the same biometric from different devices. Since the choice of biometrics is not constrained at all, we can consider this component as external to the actual core of the framework. The only thing that matters is to have an authentication channel that allows:

- *an enrolment*: an acquisition of the data that will give rise to the template;
- *a verification method*: an approach, which may be machine learning oriented, statistical or other, leading to the attribution of a score to the probe acquired following a comparison with the template;
- *a quality assessment method*: one or more approaches to assess the quality of the biometric channel in view of the acquisition and following the actual acquisition of the probe. This is fundamental in the authentication phase, but can also be extremely useful in the enrollment phase, as low quality templates can cause poor performance in the authentication phase.

The framework, potentially, can also work by implementing a single biometric, but can offer its maximum effectiveness, of course, if implemented on a multibiometric system. In the authentication phase, each biometric channel will generate an authentication state following a successful acquisition. This authentication state will contain all useful information on the acquisition, such as the score and quality level of the signal. The abstraction of biometrics lies precisely in this: the engine does not care which biometric has generated an authentication state, it is only interested in what is useful for the authentication itself.

3.3.2 Context Awareness

The context awareness is the element that will give greater sensitivity to the framework. The information it must be able to gather will depend on the context in which the security system will be contextualized. What such a component can acquire, moreover, may be useful in calculating the quality level of a specific biometric channel, but also in obtaining information that will generate, in addition to other data, an effective probe of some behavioral biometric channel. Information from the context detection can enrich the authentication state by giving the engine the possibility of making a conscious and considered choice that is not limited to a simple biometric score. The sub-components that can be implemented in context detection are:

- *Ambient detection*: collects information about the environment in which the system is installed. Information can be collected, for example, on temperature, number of subjects present, brightness, position of a target subject and so on;
- *Activity detection*: collects data on the activity of the individual being authenticated. It may indeed be useful to understand whether a user is stationary, walking or even driving;
- *Critical state detection*: on its own or even by exploiting the combined information of the previous components, it must be able to understand whether a maximum security policy or a usability-oriented policy must be preferred. A system that adapts according to a safe or critical context is able to guarantee a high number of implicit authentications (and therefore less invasiveness) and a high level of security when necessary.

3.3.3 Reliability Assessment

The reliability assessment mechanism takes care of keeping the authentication flow safe, avoiding introducing a large number of authentication states from channels that are not stable. A channel may be unstable for various reasons, and being able to detect the problem in a short time can prevent the system from raising false rejections by requesting a large number of unnecessary explicit authentications. An anomaly can be identified before this stage thanks to analyses that precedes the biometric acquisition, but also at this moment, when despite good assumptions a biometric produces varying and distant authentication scores for the same genuine subject. It is important, therefore, that in such a situation the system realizes the problem and blocks the channel. The analysis of the data from the flawed channel, however, should not be interrupted. In fact, the signal might stabilize, in which case it is necessary to start accepting again its authentication states. The reliability enriches the authentication state coming from the biometric channel with its information so that the engine can evaluate whether or not it is appropriate to consider the authentication state itself.

3.3.4 Authentication State

An Authentication State (shortened "Auth State") it's the object that encapsulates all the information mentioned so far. In fact, it could contains:

- The biometric informations, such as the score and the quality;
- Temporal information on the acquisition;
- A reference to the probe file;
- The context informations acquired at the moment of the acquisition;
- The reliability information of the channel that generated the state;
- Any information useful for the local decision controller.

3.3.5 Local Decision Controller

When an authentication state arises from a biometric channel, the decision controller processes its information in order to generate a local decision. That value will be maintained in a certain time window to allow the engine to perform its temporal analyses. The local decision controller, in fact, does not decide whether the user is authenticated or not, but triggers the analysis that will calculate the actual authentication level according to which the user will be considered authenticated or de-authenticated with respect to the current context. Moreover, depending on the goodness of the signal, an update of the enrolment template may be triggered.

3.3.6 Authentication Flow and Temporal Flow Analyzer

Authentication flow is a set of relatively recent local decisions and their authentication states that can be exploited for an analysis leading to an actual authentication level that is not calculated from a point decision, but from a set of recent states that can better define the validity of a current authentication situation. The temporal flow analyzer, on the other hand, is the component that has the task of using the information contained in the authentication flow in order to generate an authentication level between 0 and 1. Is defined as temporal because for its analysis it must take into account how recent the authentication states are.

3.3.7 Authentication Level

This is the value representing the authentication level of the user resulting from the temporal analysis. A low authentication level indicates that the system has low trust in the user who is in the authentication stream, while a high level indicates trust in its genuineness. Thanks to the context information, moreover, the evaluation of the authentication level is not flattened into a simple binary distinction, but rather the level obtained can be evaluated according to the context in order to decide whether there are conditions to maintain the authentication or fortify it with an explicit authentication. In this way, as already clarified, a low invasiveness of the system is pursued due to the high frequency of implicit authentications; moreover, is also pursued a high level of security in situations requiring more attention thanks to an explicit authentication that will be eventually prompted.

3.3.8 Security Params

It represents the configuration of the engine and all its components. It can allows to update, for example, thresholds, time parameters and any other configuration useful for the engine or the individual components of the framework.

Chapter 4

Validating Edison: a sample Implementation

4.1 Target architecture

In order to test the validity of the framework, it was necessary to run a trial of the approach that could have highlighted its potential and possibly its weaknesses. To proceed with experimentation, therefore, an example system was designed and implemented following the guidelines of the defined approach. The result is a possible implementation of the framework useful as a tool to perform experimental tests. The Edison approach has no implementation limitations on architecture, biometrics, or context and the first phase involved choosing an interesting target architecture on which to build the example application. It should be emphasized that the application that has been implemented and its architecture represent only an example since, as already clarified, the framework is potentially implementable on any architecture. So, as a first sample implementation, the most suitable target architecture was that of a smartphone and, therefore, it was decided to develop a mobile application in which each framework component would represent a software module. The choice of a smartphone architecture was dictated mainly by three reasons:

- Modern smartphones implement a large number of sensors; as Edison is a framework oriented towards multi-biometric authentication, it is possible to exploit these sensors in order to generate different biometric channels. Sensors include cameras, microphones, inertial sensors, position sensors and so on;

- The idea of an advanced CA system could go well with the continuous use of smartphones that, today, constantly accompany us in every activity;
- Looking at the evolution of smartphones in recent years, it is easy to assume that there will be even more sensors that can give rise to more possible biometrics.

Once the target architecture had been chosen, the different components provided by the framework were designed, which represented a guide, a basic approach that provides a certain type of components, which must then be designed at a logical level in detail for the specific system to be developed. In the course of this implementation, in detail, have been defined:

- target biometrics and related state-of-the-art approaches;
- the set of data acquired by the context detection module;
- an approach for the reliability module to detect biometric channels anomalies;
- the full structure of an authentication state;
- a *Fuzzy Controller* as the local decision controller of the decision engine;
- the structure of the authentication flow and the temporal analysis algorithm;
- decision engine configuration parameters;
- an advantageous deployment architecture.

4.2 Target biometrics

The choice of biometrics involved taking advantage of the basic sensors available on the target architecture, i.e. the smartphone. Following a census of available sensors, four biometrics were identified, three of which were physical and one behavioral. In particular, the example implementation of the Edison framework planned to exploit:

- *Physical biometry of the face*, using the internal camera;
- *Physical biometry of the voice*, using the microphone;

- *Physical biometry of the fingerprint*, using the fingerprint sensor;
- *Behavioral biometry of gait* , using inertial sensors.

Imagining the general use of a smartphone, in fact, it is understandable to foresee that these channels can obtain updates with a significant frequency. For example, with the camera, it would have been possible to capture ear biometrics. However, the face was preferred imagining that the frequency with which the face is visible to the camera is certainly greater than the frequency with which the ear is visible, which is generally visible when the user is about to answer a call. The smartphone is now a device that accompanies us in almost every moment of the day, so it is also easy to foresee that there will be several occasions when the user will be talking or walking. Finally, chosen as the secure explicit verification channel, the fingerprint is often used for unlocking the device.

With regard to the approaches for verification on individual biometrics, state-of-the-art systems using Machine Learning and Deep Learning were adopted, with the exception of the fingerprint, for which the android API was used directly. In particular, for the other implicit acquisition biometrics, we have:

- *DeepFace* as a state-of-the-art face verification system [21];
- *PaddleSpeech* as a state-of-the-art speaker verification system [22];
- *GaitAuth* as a state-of-the-art system for gait verification [23].

The acquisition frequency was configured for the different biometrics to pursue maximum speed (or continuity) but also taking into account the performance of individual biometrics that could take time to acquire, such as voice and gait. In particular:

- A face capture attempt takes place every second. When the heuristics or the quality signal of the biometric channel allow an authentication request, the acquisition takes place. Following a response that generates the authentication state, the next acquisition takes place;

- A voice recording on which to perform authentication lasts a few seconds. Again, the acquisition of the next probe is triggered by an authentication response of the current one;
- Gait is only acquired when a walking session has been detected by the context detection module and lasts a few seconds; if the context detection module doesn't notify a change in detected activity, the gait verification continues;
- Fingerprint is recorded both with spontaneous authentications and with authentications requested by the system at an unpredictable frequency.

4.3 Full System Architecture

The system could have been implemented entirely on a mobile application, but in order to have more convenience in the implementation of the state-of-the-art approaches for the biometrics, it was decided to divide the architecture into a frontend realized in a mobile app and a backend for the management of the verification logic of the individual biometric channels. In particular, the main framework logic was implemented on a native Android application using Java [24], while the support calls to the authentication libraries of the three biometrics channels were made available on a backend written in Python [25]. Fig. 4.1 shows the system architecture described.

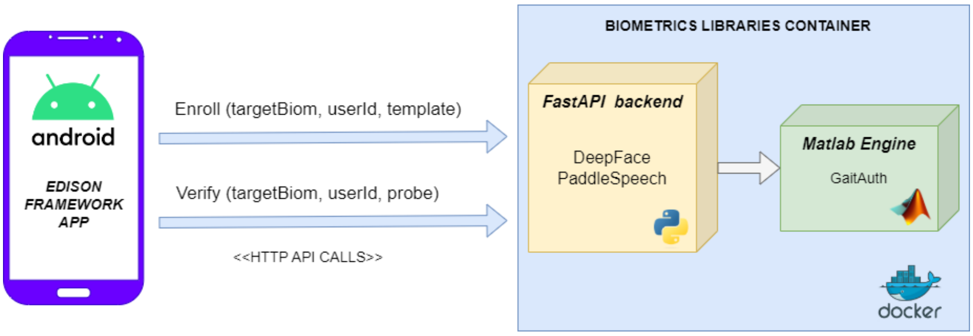

FIGURE 4.1: The architecture of Edison's example implementation.

In detail, the biometrics of face and voice were implemented directly within the backend, so once again in Python, while the gait authentication was handled through scripts executed on a pending instance of Matlab [26] with which the backend can interact. The entire backend system, consisting of the Python application and the Matlab engine, was deployed using a Docker container. The system stores templates, probes, session data and logs for each session. Furthermore, it is made known that having to test the potential of the framework regarding only the approach no security mechanism was provided for the app, the backend or the http communication channel.

4.4 The Context Detection module

The context detection module has been designed to collect information that can make the system sensitive to the conditions in which the user finds himself, what he is doing and the level of security with which he needs to be authenticated. In particular, the system has been designed to collect the following information from each sub-module:

- Ambient light and the state of the device display by the ambient detection sub-module; this information can be useful to understand, by means of heuristics, whether the user is using the smartphone, whether he has the phone in his pocket and whether he is in an indoor or outdoor environment;
- The activity of the user by the activity detection sub-module; this information can tell the system, for instance, whether the user is stationary, walking or even driving and is obtained using *Google's Activity Recognition API* [27];
- The apps used during the entire authentication flow by the critical state detection sub-module. In view of this implementation and given the target architecture, the criticality of an authentication is given by the level of security that the user declares for specific applications before the start of the authentication session. Since the sample system is contextualized on a smartphone, we can imagine that a user indicates as a critical state the moment when he is carrying out a banking transaction (thus while using his bank's app) while he may prefer greater flexibility of the system at the moment when he is simply reading an online newspaper (through a specific app).

The context detection, in the developed system, also intervenes directly in aiding the various biometric tasks by exploiting heuristics. Indeed:

- To assess the possibility of face capture, the system checks whether the display is on and whether there is sufficient brightness to assume that the smartphone is not in the pocket. In the event that DeepFace fails to detect a face in the captured photo, at the limit, the authentication state will not be entered into the authentication flow. Furthermore, a sharpness index of the photo was used as a quality value to weight the score: the higher the sharpness detected, the lower the quality of the capture;

- In order to assess the quality of the voice acquisition signal, the ambient value was used to generate a heuristic: two threshold values were experimentally defined on the basis of which the system deduces whether the smartphone is in a pocket, in an indoor environment or in an outdoor environment. If the smartphone is in a pocket or in an outdoor environment, in fact, the signal coming from the microphone is likely to have a lower quality. This result is then combined with a value representing the compatibility of the user's activity. In fact, the voice signal is more likely to be less clean if the user is running rather than standing still; similarly, it may not make sense to acquire an audio track if the user is cycling (another state detected by the context detection module). Again, at the limit, if PaddleSpeech does not detect a voice track in the actual recording, the authentication state is ignored;

- In order to acquire the inertial data, the context detection module is expected to signal an ongoing gait session. When a gait session is detected, its quality is determined based on how sure the context detection module is of an actual walking session, as what might be detected, in reality, is an erroneous running activity.

4.5 The Reliability module

The reliability module, as previously clarified, is responsible for monitoring each biometric channel in order to detect any anomalies that could interfere with the correct flow of authentication in order to block the suspect channel as long as the anomaly persists. The framework provides an anomaly detection system that is not dependent on the type of biometric: the following analysis, in fact, can be applied irrespective of the implemented biometric as it is based on the main elements of an authentication state common to any biometric. For the developed system, in detail, the chosen anomaly analysis is the calculation at each authentication state of a *Coefficient of Variation* based on the standard deviation of the decision scores in a time windowed queue. In order to obtain a normalized value, this value is divided by the arithmetic mean and the result is multiplied by one hundred in order to obtain a percentage value. The characteristics of the anomaly detection queue *q* are as follows:

- is specific to the individual biometric channel *BIOn*;
- is updated with the *enqueue* operation when a biometric channel transmits a new possible authentication state;
- has a dimension *dn* defined by the minimum number of elements to permit the analysis;
- has a time dimension *dt* that identifies how old a time *t* of an authentication state *i* can be before undergoing the *dequeue* operation;
- has a threshold *th* that represents the limit of variability a channel can have before being considered abnormal;
- has the two dimensions *dn* and *dt* and the threshold *th* configurable;
- allows the calculation of the coefficient of variation *vc* which has to be compared with the threshold *th*;
- when a referred channel is blocked because considered abnormal, it continues to accept elements in order to evaluate the unblocking of the channel.

The operation of the anomaly detection system, in detail, is algorithmically described by the diagram in Fig. 4.2. In particular, analyzing the execution flow for a generic *BIOn* biometric, a new local decision value is expected to be obtained from the individual biometric. This value is added to the queue *q* specific to the biometric under consideration with an *enqueue* operation. If the queue has a minimum size *dn*, the coefficient of variation *vc* is computed. If the computed value is below the limit *th*, then the biometric channel *BIOn* is considered safe, otherwise the received state is ignored as it is considered to be generated from an anomalous channel. Furthermore, during the entire session, queue states older than time *dt* are progressively removed. Thanks to the algorithm described, the anomaly detection system:

- performs continuous analysis on each biometric channel;
- has relatively up-to-date information on which to conduct its analysis;
- does not permanently block a channel, but only as long as it is necessary;
- does not start its analysis prematurely, but only when there are sufficient elements to conduct it.

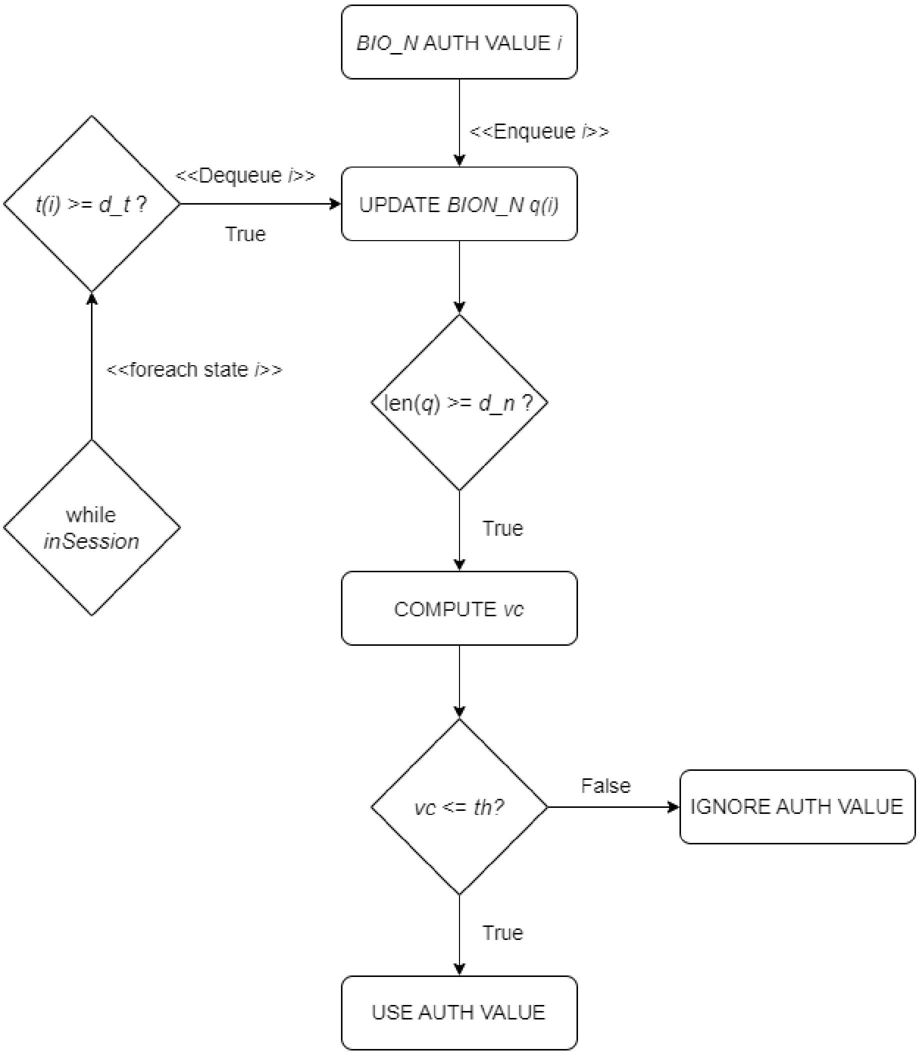

FIGURE 4.2: The anomaly detection algorithm.

Finally, the anomaly detection system takes care of checking that there is at least one channel available for continuous implicit authentication. Otherwise, at a configurable precise interval, it requests explicit authentication on a chosen secure channel (in this case, the fingerprint channel) in order to guarantee the authenticity of a user in the absence of implicit authentication conditions.

4.6 The Authentication State structure

The authentication state takes care of maintaining a snapshot of the authentication system at a certain point in time, i.e. when a new probe is acquired from a biometric channel. In fact, it contains biometric main information such as score and quality, local decision and the level of authentication that has been generated, but it also contains context and anomaly detection information.

The system stores a file containing the entire authentication flow of a session, *i.e.,* a list of authentication states. This file also contains the system configuration for that specific session, a start and end timestamp and a user identifier. By saving this file, it is possible to know all the details of an entire session, including how it evolved over time. In particular, Table 4.1 shows a dictionary listing all the elements of an authentication state.

Field	Type	Description
channel	String	The channel that generated the auth state
time	long	The auth state generation timestamp
score	float	The authentication score of the probe computed by the biometric specific verification algorithm
signalQuality	float	The computed quality of the probe
decision	float	The local decision taken by the fuzzy controller
authLevel	float	The computed auth level after the auth state
criticalState	boolean	Informs if the current auth state was generated in a critical context
probeId	long	A reference to the probe file
causedExplicitAuth	boolean	Notify if the acquired probe caused an explicit authentication for the low auth level gained with the current threshold
causedNewEnrollment	boolean	Notify if the acquired probe became a template
channelVariance	int	The channel variance computed by the anomaly detection module after this auth state
accepted	boolean	Informs if the current auth state has been accepted or not regarding anomaly detection
detectedActivity	String	The user's activity detected by the context detection module
devicePlace	String	Through a lighting heuristic, it indicates whether the device was in pocket or in an indoor/outdoor environment
refusedCauseCode	int	An error code that informs why the auth state is not valid
refusedCauseText	String	Textual explanation of the refusedCauseCode
totalExplicitAuths	int	A count of explicit authentications obtained

TABLE 4.1: The Authentication State structure.

4.7 Local Decisions: Fuzzy Controller Design

As local decision controller a *Fuzzy Controller* has been implemented. A Fuzzy Controller is based, as its name suggests, on fuzzy logic. It is a system that analyzes analog input values in terms of logical variables that take on continuous values between 0 and 1, in contrast to classical or digital logic, which operates on discrete values of 0 and 1. Using a fuzzy controller, it is possible to obtain a local decision value from several input values instead of obtaining a state that discreetly indicates authentication or non-authentication. A generic fuzzy controller involves the following steps:

- the fuzzification of the controller inputs;
- the execution of the rules of the controller;
- the defuzzification of the output to a crisp value.

In the example system, a Mamdani model was implemented with two inputs: *auth score* and *signal quality*. This controller, in fact, takes care of weighting the score returned by a biometric by exploiting its signal quality. Finally, the output of the combination of the two values is obtained: the *trustness*. Once identified, the controller inputs and outputs were modeled as logical variables through experimentally defined triangular membership functions . The distributions of the auth score, signal quality and trustness variables are shown in Fig. 4.3, Fig. 4.4, and Fig. 4.5, respectively.

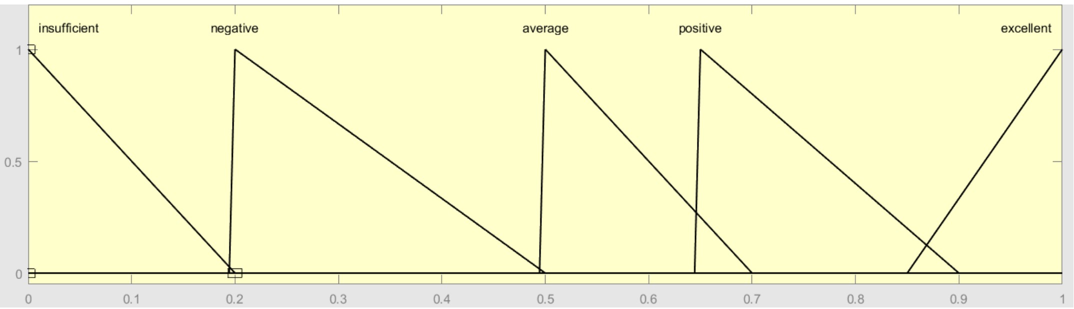

FIGURE 4.3: Definition of the auth score variable with the values: insufficient, negative, average, positive and excellent.

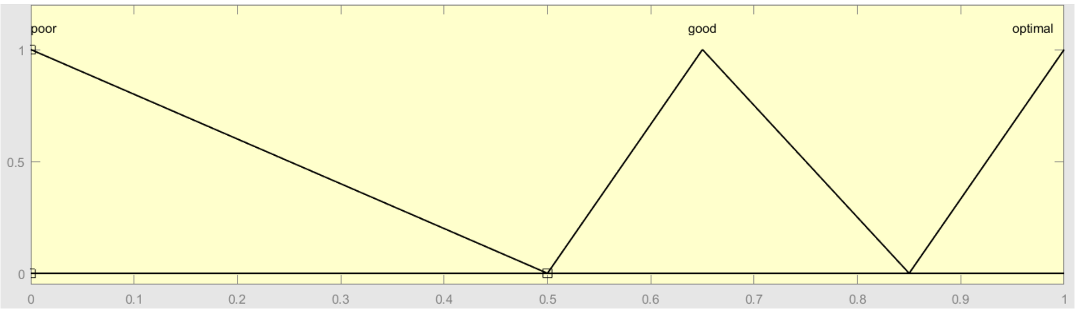

FIGURE 4.4: Definition of the signal quality variable with the values: poor, good and optimal.

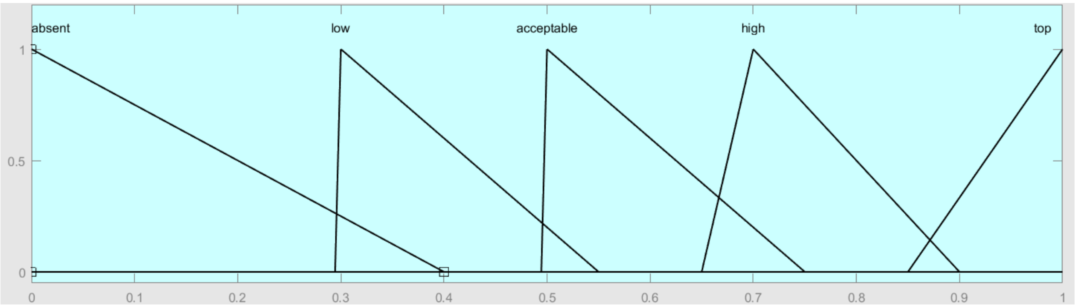

FIGURE 4.5: Definition of the trustness variable with the values: absent, low, acceptable, high and top.

Once the logical variables had been defined, the rules according to which the controller need to operate in order to combine the two values were defined. Specifically, the rules are as follows:

1. *If (Auth_Score is excellent) and (Signal_Quality is optimal) then (Trustness is top)*
2. *If (Auth_Score is positive) and (Signal_Quality is optimal) then (Trustness is high)*
3. *If (Auth_Score is average) and (Signal_Quality is optimal) then (Trustness is acceptable)*
4. *If (Auth_Score is negative) and (Signal_Quality is optimal) then (Trustness is low)*
5. *If (Auth_Score is insufficient) then (Trustness is absent)*
6. *If (Auth_Score is excellent) and (Signal_Quality is good) then (Trustness is high)*
7. *If (Auth_Score is excellent) and (Signal_Quality is poor) then (Trustness is acceptable)*
8. *If (Auth_Score is negative) then (Trustness is low)*
9. *If (Auth_Score is positive) and (Signal_Quality is good) then (Trustness is high)*

The application of these rules allows, in a concrete example shown in Fig. 4.6, to lower the trustness value if the biometric score was positive but the signal quality was relatively poor. It is emphasized that the centroid method was chosen as the defuzzification method.

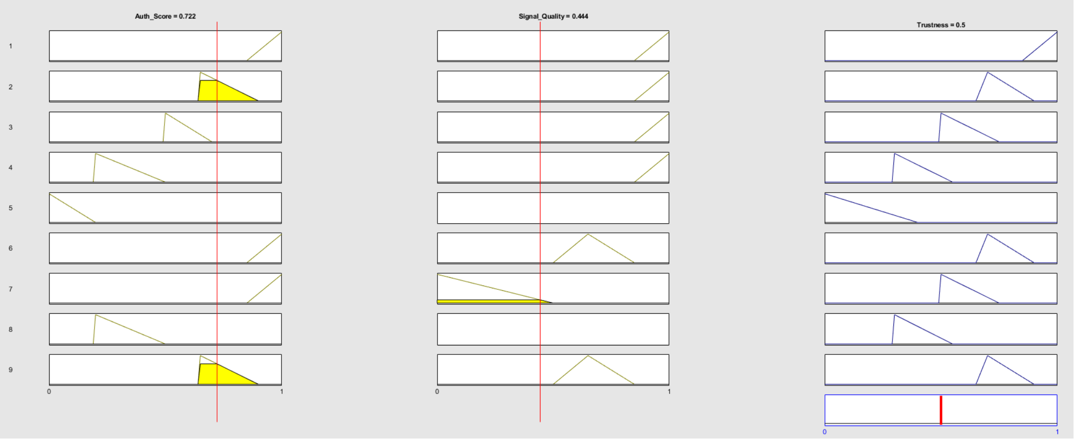

FIGURE 4.6: Testing the rules: an auth score of ≈ 0.7 and a signal quality of ≈ 0.4 returned a trustness value of 0.5.

Finally, in Fig. 4.7, the entire distribution space of trustness with respect to auth score and quality can be observed.

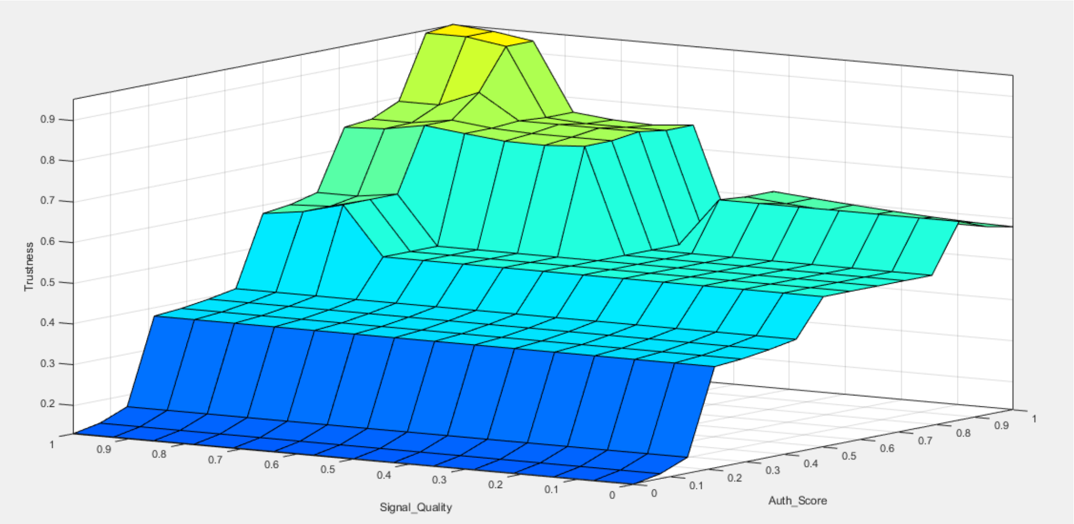

FIGURE 4.7: Plot of the trustness distribution on the auth score and quality inputs.

4.8 Flow Analyzer: a Temporal-oriented Weighted Mean

As a temporal flow analyzer we implemented a *Temporal-oriented Weighted Mean* that operates on the authentication flow, that represents the dataset on which the decision analysis will be performed to obtain the authentication level. The authentication flow, to allow the computation of the weighted mean, must be a queue with these characteristics:

- is *unique* for all biometric channels. The moment an authentication state enters the authentication flow, the channel of origin is no longer important: all that is needed for analysis purposes is to maintain in the flow the main information of the valid authentication state. In this case, the analysis is performed on the local decisions and the acquisition times of each state in the queue to compute the authentication level;
- is updated with the *enqueue* operation when a biometric channel transmits an authentication state that has passed the check of the anomaly detection module;
- has a configurable time dimension *dt* that identifies how old a time *t* of a state *i* may be before it undergoes the *dequeue* operation; it is plausible to assume that authentication states that are too old should no longer be subject to analysis;
- can accept implicit as well as explicit authentication states;
- provides a *distribution of temporal weights* among its elements in order to allow an analysis weighted by temporal factors. It is plausible to assume, in fact, that the more recent elements should have a higher weight than the less recent ones.

The key element of the described queue that allow the temporal-oriented mean computation are exactly the temporal weights. These weights are a homogeneous and ordered distribution of values that have been arranged in three categories as shown in Fig. 4.8. In particular:

- Older values, labeled *OLD*, will range from a minimum weight of 0.1 to a maximum weight of 0.4;
- Recent values, labeled as *RECENT*, will range from a minimum weight of 0.5 to a maximum weight of 0.7;
- The extremely recent values, labeled *ACTUAL*, will range from a minimum weight of 0.8 to a maximum weight of 1.

From the configuration parameter *dt*, the queue and the weights distribution are automatically configured. The parameter indicates the time size of the queue and is expressed in seconds. From this value, the three previously defined time partitions are obtained. Subsequently, the sub-partitions on which the weights of the respective category can be equidistributed are calculated.

When an authentication state is inserted into the authentication flow, the weight attributed to the state itself is the maximum one, since it's a state that has just been received. Moreover, the weights of the other states that are still in the queue are recalculated by choosing a value that depends on the main logical position of the elements in the temporal queue (OLD, RECENT, or ACTUAL) and their even more precise position in the temporal sub-partition.

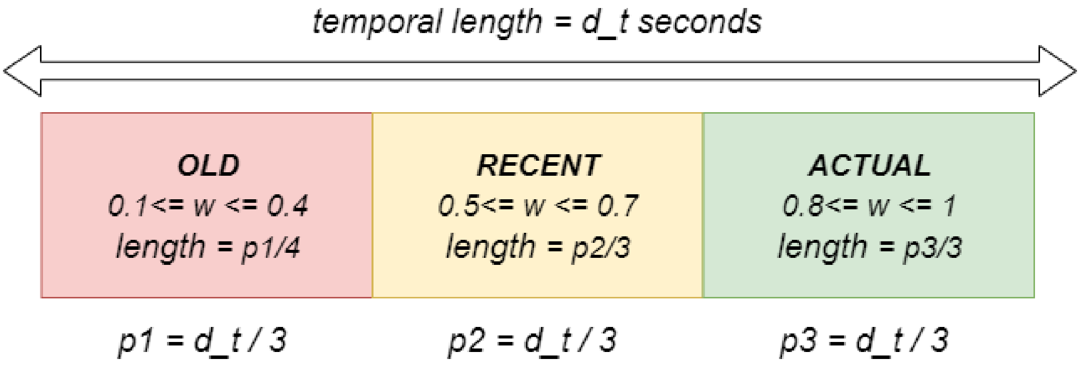

FIGURE 4.8: Distribution of weights on the time queue of configurable size *dt*.

The temporal flow analyser, at this point, has everything it needs to be able to generate the authentication level. The computation, in fact, is done by calculating a time-weighted average of the authentication states present in the authentication flow. In this computation, the recent states will have a greater weight than the oldest states, which will still retain some importance (even if minimal) if they still are in the queue. Based on the level of criticality provided by the critical context detection module, the authentication level obtained will be compared with a certain threshold. Thresholds are configurable according to the level of criticality within the application.

4.9 Using the Mobile App

The mobile app has been equipped with a graphical interface to allow the system state to be kept under control with key information and to manage the necessary interactions. In particular, Fig. 4.9 shows the main graphical interface of the mobile application in an authentication session. In particular, three buttons are available in the upper part of the app: the first one allows to enter the URL of the backend (to allow different deployments of the backend in the various development and experimentation phases), the second one allows to choose the critical apps from a list that gathers all the apps installed in the device (Fig. 4.10), while the third one (Fig. 4.11) allows to carry out an enrolment of a single biometric or a sequential enrolment of all the biometrics with also a reset of the participant ID. Next, after the current authentication level label, there is a line chart that allows one to observe the local decisions recently obtained with the local decision and the source channel (with a click on the decision point); there are also two dotted lines showing the current authentication level and threshold. For the authentication flow, it is possible to observe how many elements are currently in the queue, how many authentication states have been processed during the entire session and, finally, the total number of explicit authentications that have taken place, counting both those requested by the system and those spontaneous. In the central part of the layout, on the other hand, up-to-date informations on the various biometric channels are shown. In particular, for each implicit channel, the following fields are shown: the latest decision and scores, the latest quality levels, the current variation coefficient and a count of rejected states (due to anomalies) out of the total states received. Finally, in the last part of the layout, there are some informations from the context detection module; in particular, the following fields can be displayed: the lighting level, the current detected activity with a confidence level in percentages and the name of the last critical app used by the user. The system, moreover, allows the user to manage certain configuration parameters. In particular, as shown in Fig. 4.12, the app allows parameters to be set for both the authentication flow and the anomaly detection module. In particular, for the authentication flow it is possible to specify the thresholds for the basic and critical authentication context, the temporal dimension of the authentication flow on which the analysis for the authentication level will take place and the threshold which allows a probe with a certain quality and score to be acquired as a new template. For anomaly detection, on the other hand, it is possible to specify the minimum size that the queues must respect in order to begin analysis, the maximum time a state can remain in the queue, the threshold for the maximum variance allowed before blocking the channel and a scan interval in seconds that marks the time in which, in the absence of non-anomalous channels, explicit authentication must be requested.

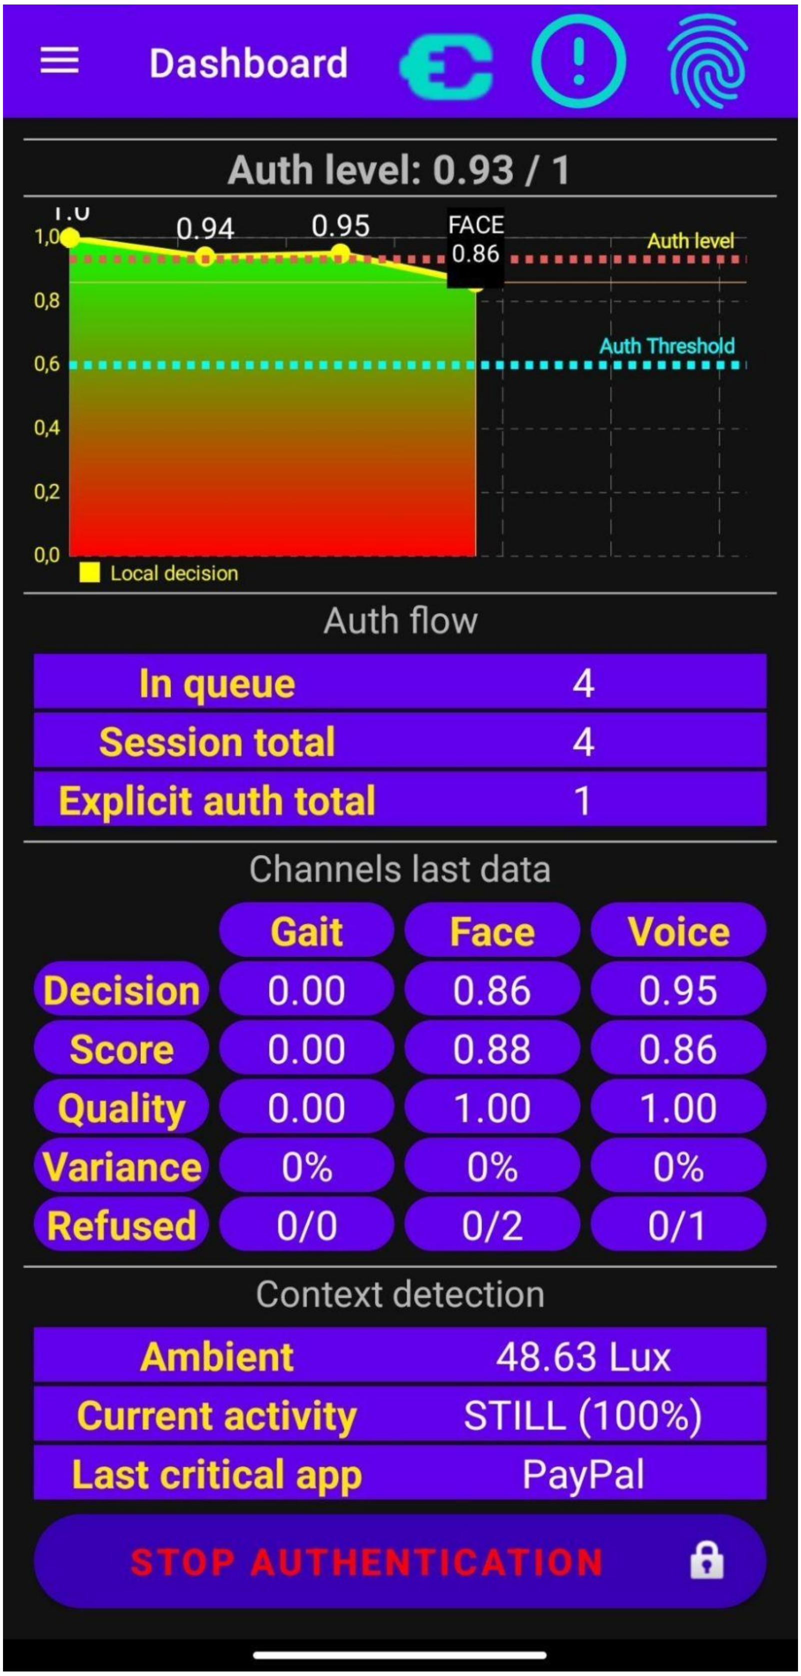

Dashboard

Auth level: 0.93 / 1

Auth flow

In queue	4
Session total	4
Explicit auth total	1

Channels last data

	Gait	Face	Voice
Decision	0.00	0.86	0.95
Score	0.00	0.88	0.86
Quality	0.00	1.00	1.00
Variance	0%	0%	0%
Refused	0/0	0/2	0/1

Context detection

Ambient	48.63 Lux
Current activity	STILL (100%)
Last critical app	PayPal

FIGURE 4.9: The application dashboard during a session.

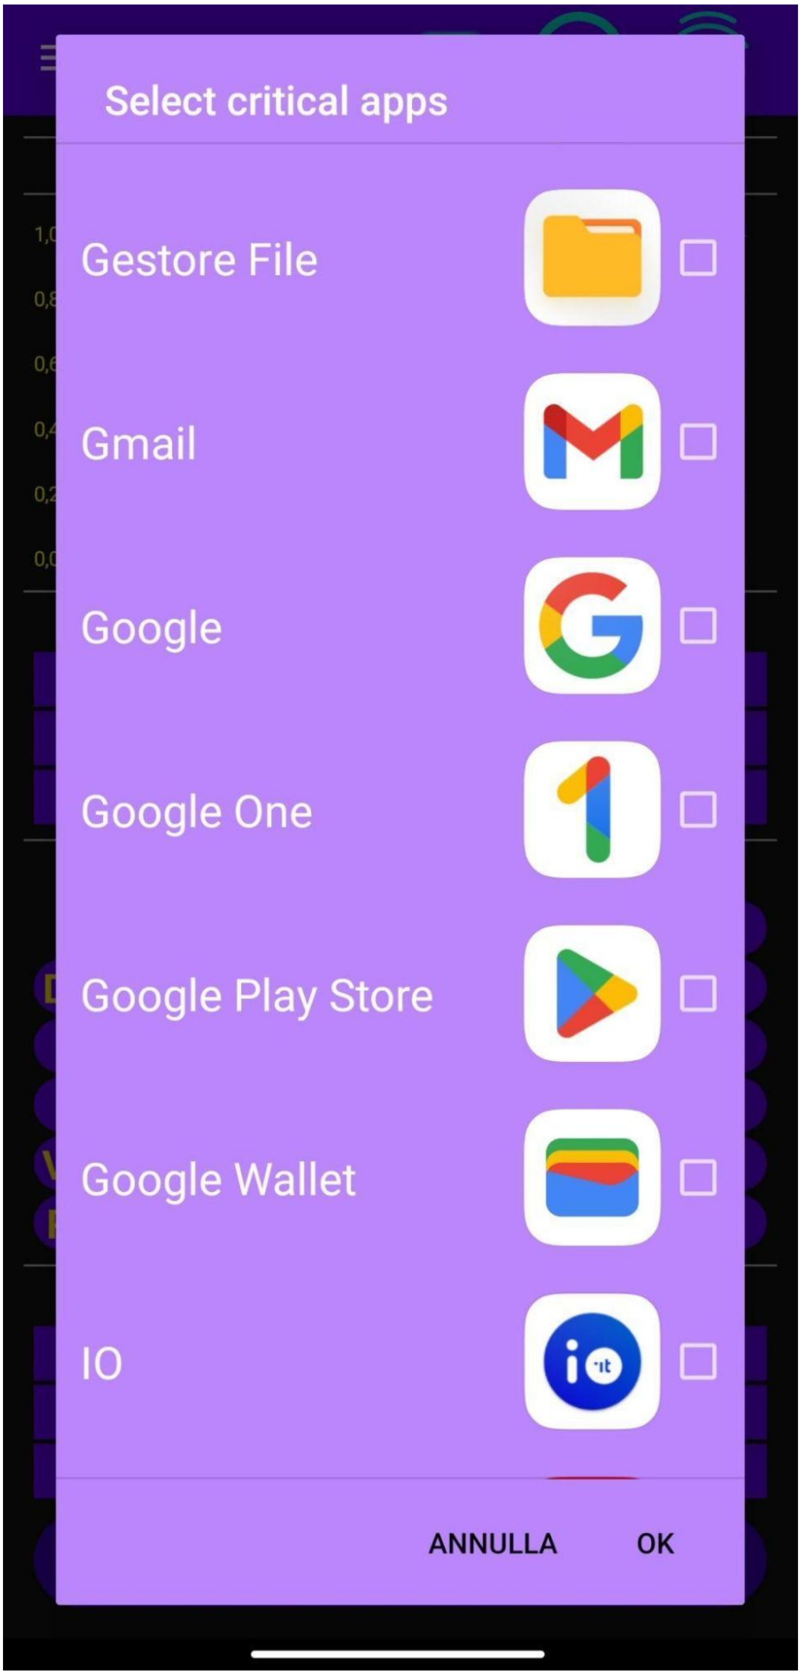

Select critical apps

Gestore File □

Gmail □

Google □

Google One □

Google Play Store □

Google Wallet □

IO □

FIGURE 4.10: The critical apps selection.

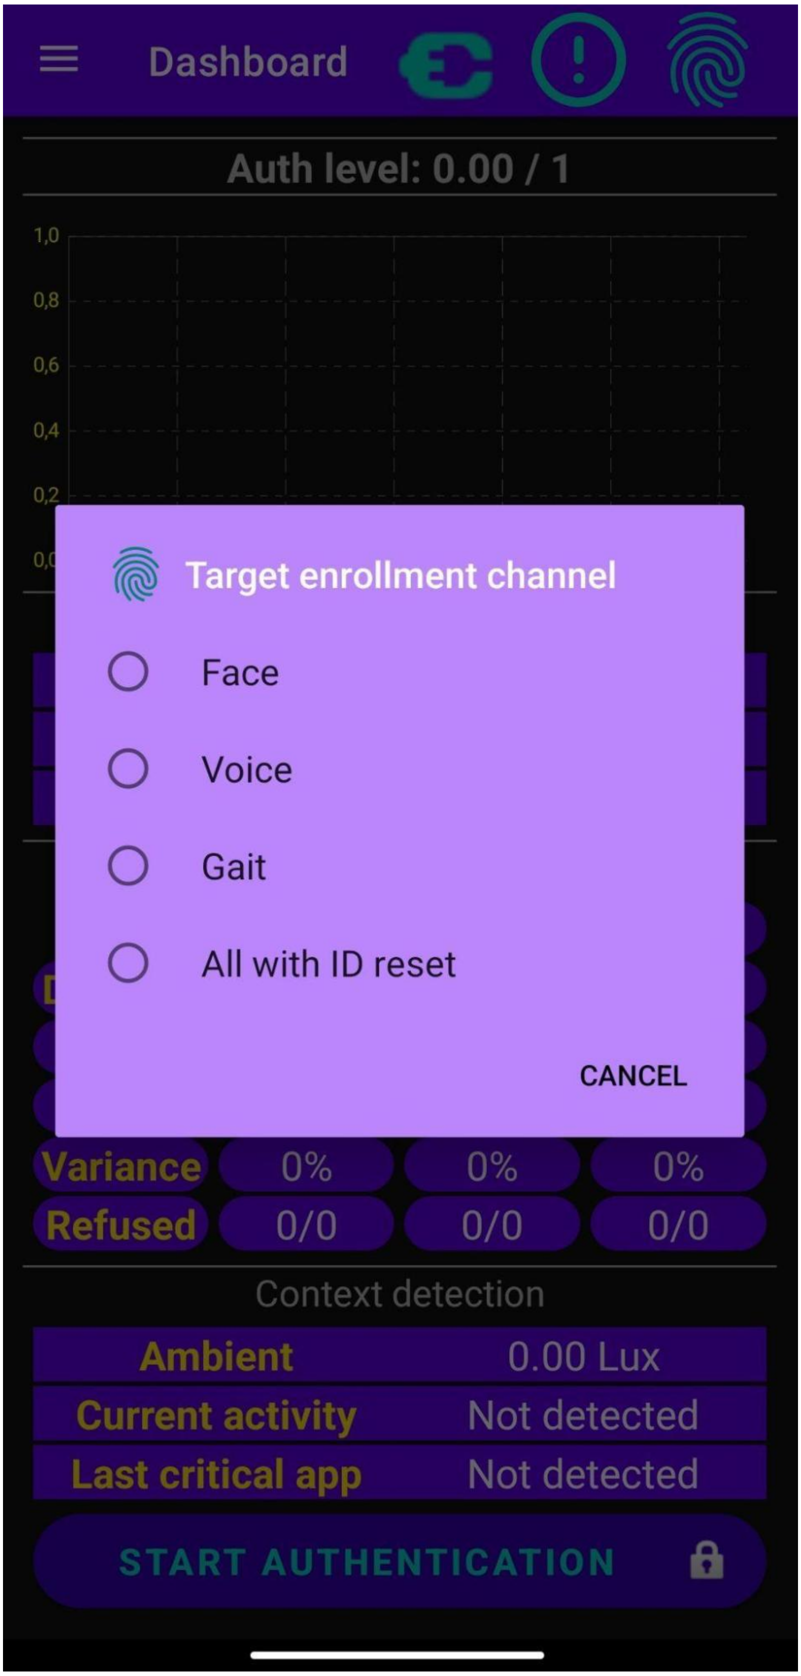

Target enrollment channel

○ Face

○ Voice

○ Gait

○ All with ID reset

CANCEL

Variance	0%	0%	0%
Refused	0/0	0/0	0/0

Ambient	0.00 Lux
Current activity	Not detected
Last critical app	Not detected

FIGURE 4.11: The enrollment management.

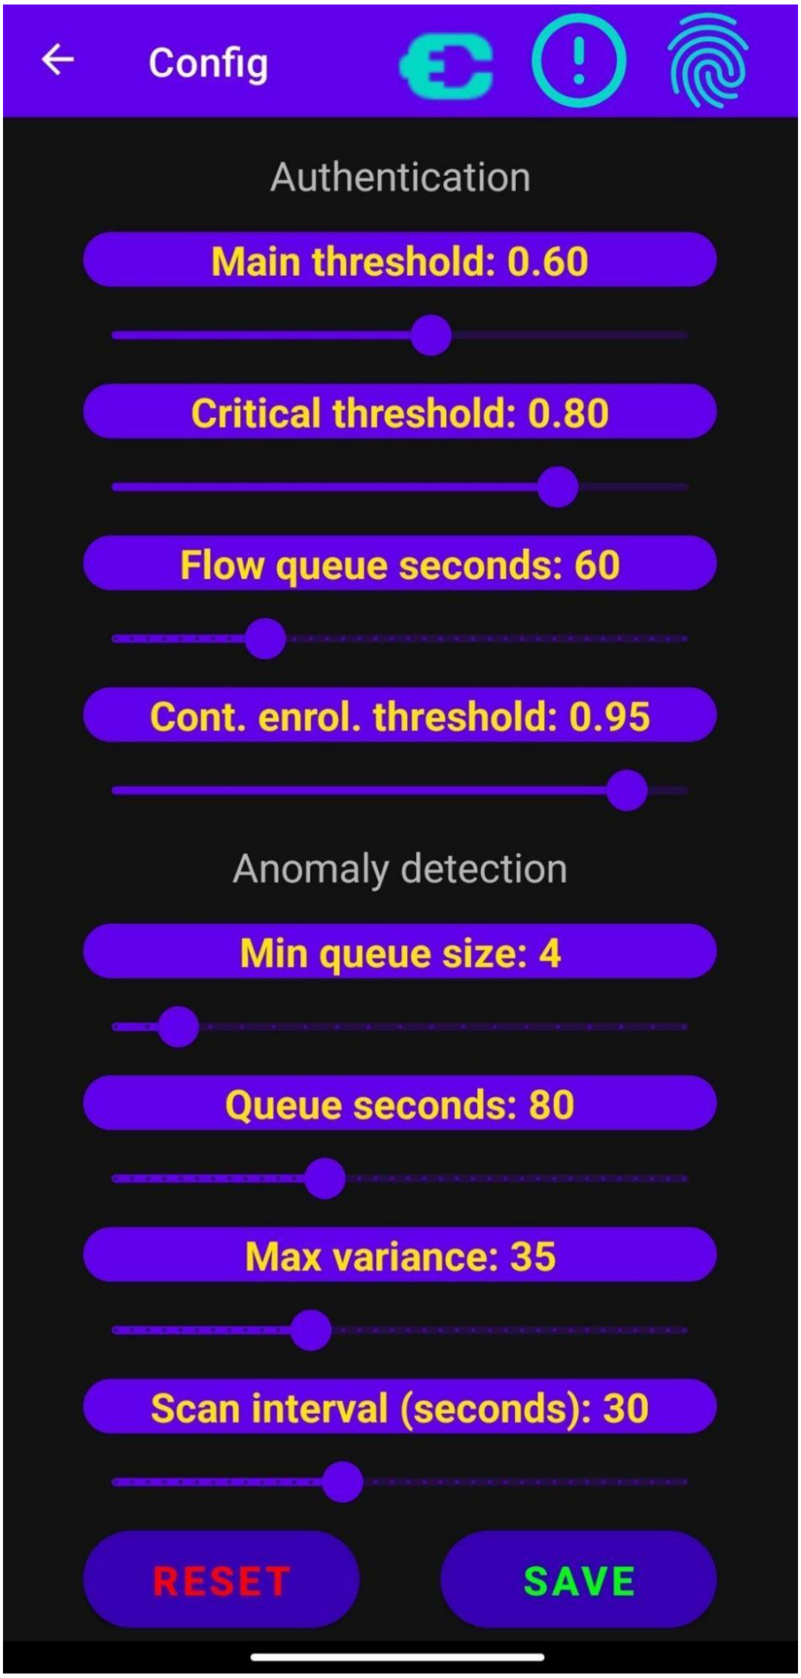

Authentication

Main threshold: 0.60

Critical threshold: 0.80

Flow queue seconds: 60

Cont. enrol. threshold: 0.95

Anomaly detection

Min queue size: 4

Queue seconds: 80

Max variance: 35

Scan interval (seconds): 30

FIGURE 4.12: The parameters configuration page.

Chapter 5

Validating Edison: Experiments and Results

5.1 Experimental Protocol

The experimentation included an authentication session for each participant using a smartphone running Android 12 with the mobile app installed. Specifically:

- there were 12 participants;
- each session included an enrollment of the individual participant;
- each session, following enrollment, included 3 tests, each lasting a few minutes;
- it was possible to use the three implicit authentication biometrics implemented: face, voice, gait.

The participant clicked on "Start authentication" button to start a specific session test. Explicit authentication via fingerprint has been requested at each start. For both initialization explicit authentication and any other explicit authentication required, the participant selected "Use Pin" and entered “0000”. This action replaced explicit authentication over a secure channel, which in the developed mobile app is the fingerprint. This choice allowed the participant to proceed in a facilitated manner throughout tests. In case of de-authentications during sessions, the system required explicit authentication by locking the device. Explicit authentication, in order to easily manage the entire trial, required entering the same PIN code used to start the session. The system, moreover, required as input a simple incremental numeric user ID that has been communicated in advance. This ID is uniquely associated with the participant. Thereafter, EDISON captured the three implemented biometrics for the enrollment. In particular, face has been captured through a photo, voice through a 5-second recording (using a random sentence of participant choice), and gait through a 20-second walking session. After the enrollment, for each of the three tests, suitable conditions were recreated to test specific features of the framework simply by telling the participant what actions to perform.

5.2 Session Execution and Extracted Metrics

Guidance has been offered to the participant throughout the session in order to keep the testing environment controlled. As mentioned, each participant's session included the execution of 3 tests designed to test the specific potential and criticality of the framework: *Effectiveness of Authentication Flow*, *Fault Tolerance* and *Invasiveness*.

5.2.1 Effectiveness of Authentication Flow

Aims to demonstrate the effectiveness of analysis on a time stream of authentication states versus a point decision that may be susceptible to noisy signals causing explicit authentication when not needed. The test included the following steps:

- The anomaly detection system has been configured by assigning a maximum allowable variance to a relatively high value to avoid getting channel blockage due to alternating clean and noisy values that has been voluntarily offered to the system for the current test;

- Positive acquisitions and noisy captures has been intermittently caused on a single channel continuing to offer positive acquisitions to other channels trying to achieve the maintenance of a legitimate authentication state (with a lower but still sufficient level of authentication) where a timely decision would have de-authenticated the user following a simple threshold comparison. For the voice biometric, intonation has be changed to introduce noisy signals. For face biometrics, on the other hand, it was possible to cause blurry image captures. Finally, for gait biometric, the possible choice was to forcibly change walking style. With such modes, intra-class variations would be achieved, but it was also possible to proceed by introducing another participant's biometrics;

- Towards the end of the test, an impostor was instantaneous brought in the flow to check how much time was required for de-authentication.

The calculated metrics for the test are the following:

- *Flow Genuine Effectiveness (FGE)*: it calculates in the authentication flow the number of explicit authentications *nExplAuths* required by the system and divides it by the number of *nLowDec* authentication states that have a lower local decision score than the threshold; basically, it involves exploiting a false rejections rate for the calculation; The higher the value, the more we can consider that the flow has been effective in maintaining authentication of the genuine versus a point decision that would have de-authenticated it;

$$FGE = (1 - \frac{nExplAuths}{nLowDec}) * 100$$

- *Impostor Reaction Time (IRT)*: it simply calculates the elapsed time taken by the system to de-authenticate the impostor from the first authenticate state obtained when it was introduced in the session. If the value obtained results in 0, the system has raised an instantaneous de-authentication.

5.2.2 Fault Tolerance

Aims to demonstrate the effectiveness of the reliability assessment by testing the anomaly detection module that can temporarily prevent the entry of authentication states into the authentication flow that have not been stable for a certain period of time for a specific channel. The steps for such a test were as follows:

- The anomaly detection system has been configured by assigning an ideally low value to a maximum allowable variance;
- A one-channel oscillation was caused by intermittently bringing in the flow an impostor instead of the genuine user. Specifically, in turn, there has been an authentication with probe belonging to the genuine user and one with probe

belonging to the impostor, all while the genuine user continued to offer legitimate authentications on other channels;

- The system has been monitored on the activation of anomaly detection module for the oscillating channel.

It is intended to test that if there were no anomalies all channels are used achieving a certain performance in terms of implicit authentication rate, while if some channels were blocked by the anomaly detector the system would remain performant compared to the previous ideal situation.

The calculated metric for the test is the *Implicit Authentication Rate (IAR)*, *i.e.,* the number of implicit authentications *successAuths* divided by the number of total authentication states *totAuths*. This value has been computed before and after the activation of the anomaly detection system to compare authentication rates before and after a channel blocking.

$$IAR = \frac{successAuths}{totAuths} * 100$$

5.2.3 Invasiveness

The test served to compare the invasiveness levels of the system in a low-security context and in a high-security context to verify that we have a low degree of invasiveness in the former case and a high degree of security in the latter. For this purpose, positive authentications has been performed on different channels and explicit authentications required under normal and critical conditions has been counted. The calculated metric for the test is the *Invasiveness Index (INVI)*: it calculates the number of explicit authentications *explAuths* required by the system and it divides it by the number of total authentication states *totAuthsNorm*. The metric has been calculated in the same session test in a critical context and in a non-critical context to compare results before and after a high level security context switch.

$$INVI = \frac{explAuths}{totAuthsNorm} * 100$$

5.3 Results

Analysis were performed using Python scripts to analyze the authentication sessions for each participant in order to calculate the indicated metrics with explanatory plots.

The results show an effectiveness of authentication flow for most participants over individual point decisions. An example demonstrating the effectiveness of flow in keeping a genuine user authenticated despite frequent noisy signals is shown in Fig. 5.1. Specifically, the participant was required to hold the smartphone offering facial and voice authentications most of the time, occasionally introducing intra-class variations by radically changing facial expressions as shown in Fig. 5.2. As we can observe, the channel with the noisy signal did not caused de-authentication. In contrast, the introduction of an impostor into the authentication stream instantaneously caused a de-authentication.

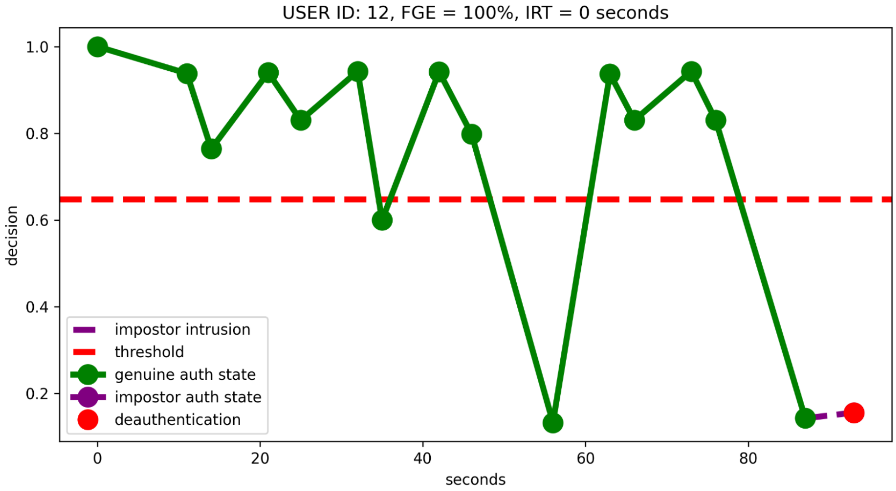

FIGURE 5.1: A good results from a flow effectiveness test.

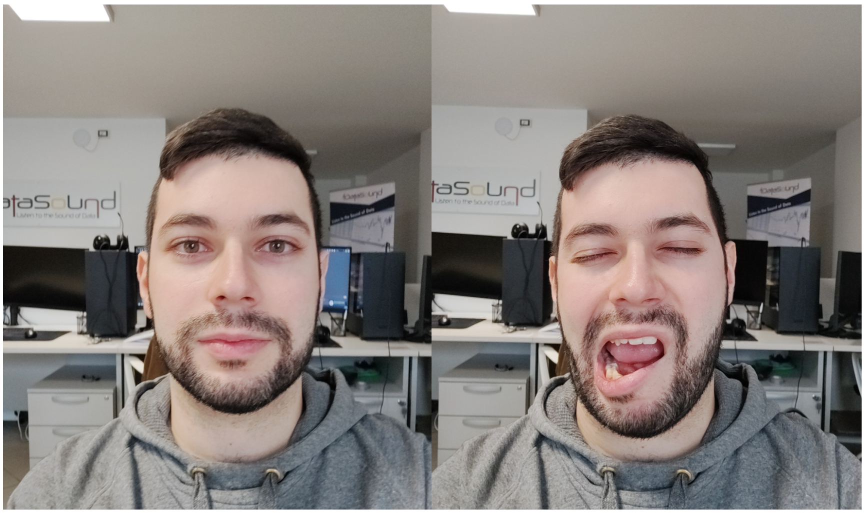

FIGURE 5.2: An intra-class variation example.

The speed of de-authentication, however, was sometimes not so performant in all cases. In fact, since the authentication flow is temporal, at relatively high recent authentication states, the impostor is detected in a longer time and with more acquisition needed, as shown in another session in Fig. 5.3. However, even if the time is high, the time between different biometrics to offer acquisitions and decisions should be considered.

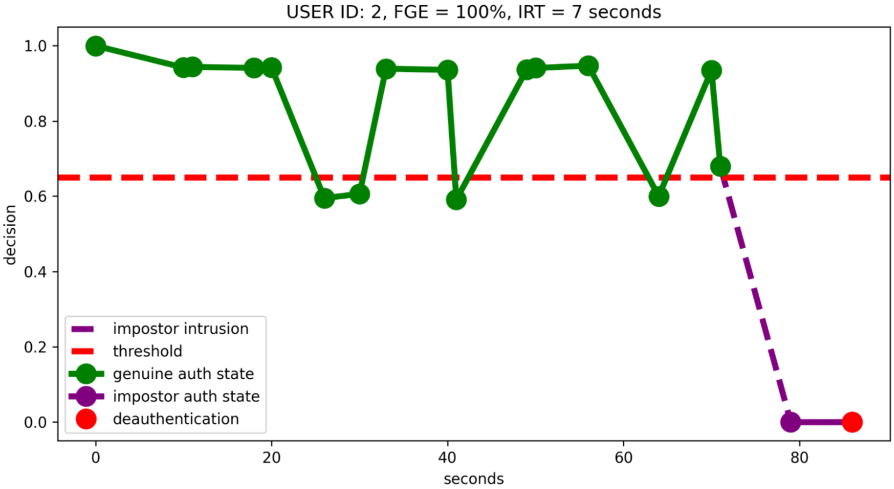

FIGURE 5.3: An example of flow effectiveness test results where more time was needed for de-authentication.

As for the fault tolerance test, despite the excellent values of the calculated metric, the results showed how an improvement of the anomaly detection system is needed. Specifically, in some cases an anomaly channel was blocked correctly, as shown in Fig. 5.4, and indeed the implicit authentication rate remained unchanged before and after the anomaly on a channel. In other cases, however, anomalies were blocked too aggressively, as shown in Fig. 5.5. In these cases, the risk is that a possible biometric channel would be blocked because it was considered anomalous when in fact it was about to signal the actual presence of an impostor on the channel. So, it becomes necessary to implement mechanism to avoid this scenario and is needed to endow anomaly detection with increased sensitivity to momentary noisy signals in order to block a channel only if it consistently exhibits effective anomalies.

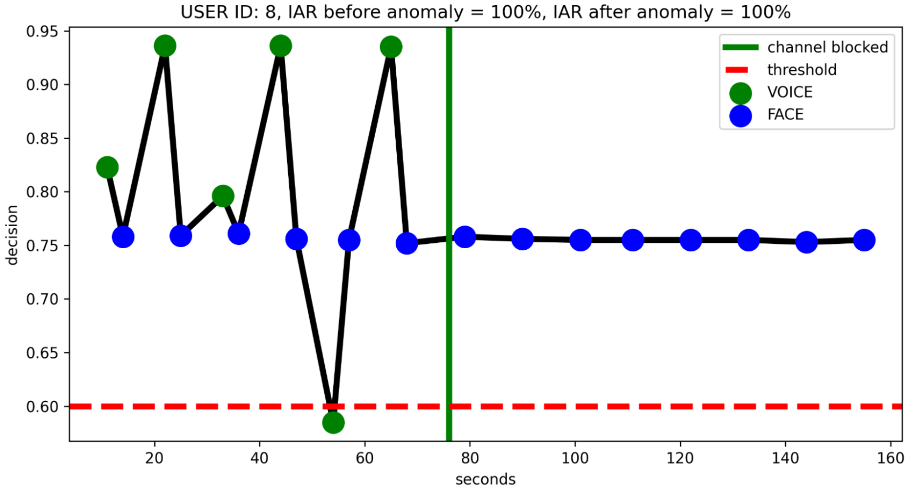

FIGURE 5.4: An effective anomaly detection on the voice channel.

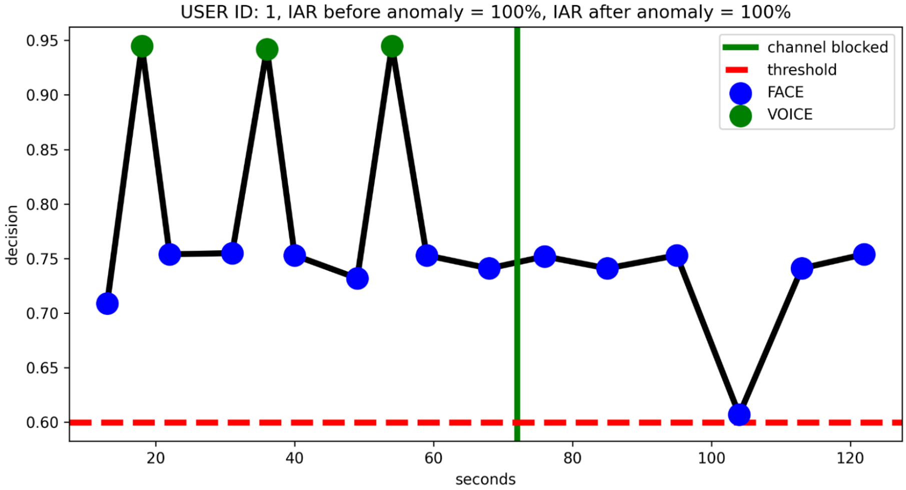

FIGURE 5.5: An aggressive deleterious anomaly detection case.

Finally, regarding the system invasiveness test, it is evident that in most cases the system was very usable under non-critical conditions, providing a high level of success implicit authentications without requiring explicit and invasive authentications from the user. In many cases, moreover, invasiveness properly increased the moment the participant entered a critical context. An example of invasiveness assessment is shown in Fig. 5.6. As we can observe, although the authentication flow is almost similar between the two contexts, the invasiveness in the critical context has increased by 22%.

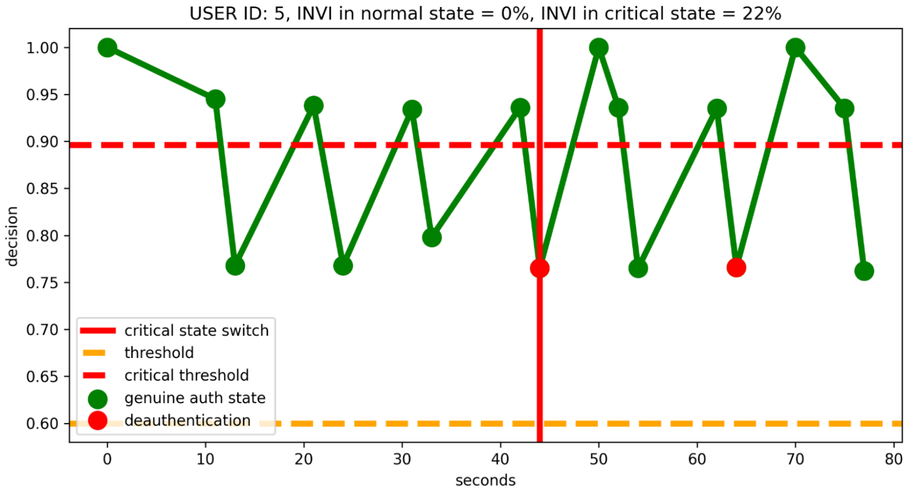

FIGURE 5.6: An example of the difference between a normal state and a critical state invasiveness.

Thanks to experimentation, we are now able to answer the research questions posed at the outset. In particular:

RQ1. Is it possible to build a CA Framework by abstracting its input and thus not being tied to specific biometrics or architectures?
At the operational level, we were able to create a working system that did not depend on specific biometrics or architectures. In fact, the framework did not effectively limit either the choice of a smartphone architecture or the choice of a particular biometric that we have chosen.

RQ2. Can a Temporal-oriented Analysis of an Authentication Flow be more effective in maintaining the authentication of a genuine user than considering a current single decision?

The results indicate that a temporal analysis performed on an authentication flow is definitely more effective than a local evaluation of an authentication state versus a threshold and that the de-authentication of an impostor user despite the authentication flow does not take an excessive amount of time.

RQ3. Can a channel Reliability Assessment mechanism be useful to maintain an high authentication rate in a CA system?

An anomaly detection system can be most useful for this purpose, but mechanisms should be provided to recognize short anomalies that should not be considered. In addition, such a mechanism should prevent aggressive evaluation of a score change that could cause deleterious blockages of a channel.

RQ4. Can a Context Awareness mechanism make a CA system less invasive in a normal condition and more secure in high-security settings?

A context awareness mechanism can actually make a system less intrusive in a context where flexibility is allowed and, at the same time, can allow more stringent requirements in a high security context.

Chapter 6

Conclusion and Future Work

Analyzing the state of the art in the context of Continuous Authentication systems, some interesting aspects emerged that seem not to have been evaluated. The generalization of an approach that did not depend on specific biometrics or architectures and particular configurations, the use of context and reliability information to enhance, weight, and enrich a simple score by one or more biometrics, and an analysis of an authentication flow with emphasis on the temporal factor are the key elements of our research work. Thanks to these key points, it's possible to move away from the concept of point decision making, of simply finding new biometrics or new approaches to process them, and to proceed to interesting insights such as a new way of evaluating biometrics using temporal, context and reliability factors and, consequently, of moving away from the dual concept of authenticated/unauthenticated and moving to a higher level, sounding out a new way of looking at acquired data in order to evaluate a maintenance of an authenticated state and, eventually, to arrive at a loss of that state, i.e., de-authentication.

In the course of this thesis work, these outlooks resulted in Edison, a novel Generalizable, Reliable, Context-sensitive and Temporal-oriented Continuous Authentication Framework. The approach of the framework we proposed has been submitted to an initial validation through an experimentation. In view of the experiment, a possible implementation of the framework was designed and subsequently developed through a system composed of an Android mobile application and a backend capable of offering support for the different biometric channels of implicit authentication chosen, namely face, voice and gait. Thanks to the developed sample system, it was possible to run a trial of the approach on twelve participants, and several metrics were defined and calculated that could help to highlight potentials and limitations of the proposed framework.

The trial of the system involved 3 tests for each participant. The first tested the effectiveness of an authentication flow against a point decision and the responsiveness of the system to an impostor; the second test, instead, allowed the reliability assessment evaluation and, finally, the third test checked the invasiveness of the system based on a switch between a normal context and a critical context.

Experimentation has obtained positive results and allowed us to answer the research questions we had been asking ourselves: the authentication flow can be very useful compared a point decision to maintain a genuine user authenticated with non invasive implicit authentications, a reliability system using an anomaly detection mechanism can maintain optimal performance and ensuring only clear channel data and a context detection can results in a low invasiveness and at the same time a high level of security when needed. Moreover, with a framework as designed, it's actually possible to achieve a generalizable approach that is not tied to specific biometrics or architectures, because we were able to achieve an effective concrete implementation without having design-related impediments from the approach itself.

On the other and, however, the experimentation also shown that it is necessary to equip anomaly detection with more intelligent mechanisms, which, ideally, should be able to identify in the time stream whether a channel has oscillations of a very short instant in such a way that the channel itself is not considered anomalous. In particular, a neural network could be effective to perform anomaly detection with dynamic periodic training based on the acquired time series by exploiting, for example, a buffer that keeps track of positive and negative examples to do retraining. In addition, some clever mechanism should be introduced to recognize collapses in decision scores that blatantly indicate the presence of an impostor and not a simple noisy acquisition. As last eventual future development, moreover, could involve automatic testing of biometrics following enrollment, both for the purpose of checking their quality and for intelligent automatic setting of certain framework parameters in order to have adaptability of the approach on the specific final user.

Bibliography

[1] S. Ayeswarya and Jasmine Norman. “A survey on different continuous authentication systems”. In: *International Journal of Biometrics* 11.1 (2019), p. 67. DOI: 10.1504/ijbm.2019.096574. URL: https://doi.org/10.1504/ijbm.2019.096574.

[2] Xiaofeng Lu et al. “Continuous authentication by free-text keystroke based on CNN and RNN”. In: *Computers and Security* 96 (Sept. 2020), p. 101861. DOI: 10.1016/j.cose.2020.101861. URL: https://doi.org/10.1016/j.cose.2020.101861.

[3] Nyle Siddiqui, Rushit Dave, and Naeem Seliya. *Continuous Authentication Using Mouse Movements, Machine Learning, and Minecraft*. 2021. DOI: 10.48550/ARXIV.2110.11080. URL: https://arxiv.org/abs/2110.11080.

[4] Mario Frank et al. “Touchalytics: On the Applicability of Touchscreen Input as a Behavioral Biometric for Continuous Authentication”. In: *IEEE Transactions on Information Forensics and Security* 8.1 (Jan. 2013), pp. 136–148. DOI: 10.1109/tifs.2012.2225048. URL: https://doi.org/10.1109/tifs.2012.2225048.

[5] Hataichanok Saevanee et al. “Continuous user authentication using multimodal biometrics”. In: *Computers and Security* 53 (Sept. 2015), pp. 234–246. DOI: 10.1016/j.cose.2015.06.001. URL: https://doi.org/10.1016/j.cose.2015.06.001.

[6] Matthias Trojahn and Frank Ortmeier. “KeyGait Framework for Continuously Biometric Authentication during Usage of a Smartphone”. In: *MOBILITY 2013, The Third International Conference on Mobile Services, Resources, and Users*. Jan. 1, 2013. published.

[7] Mohammed Abuhamad et al. *Sensor-based Continuous Authentication of Smartphones' Users Using Behavioral Biometrics: A Contemporary Survey*. 2020. DOI: 10.48550/ARXIV.2001.08578. URL: https://arxiv.org/abs/2001.08578.

[8] Abbas Acar et al. "WACA: Wearable-Assisted Continuous Authentication". In: *2018 IEEE Security and Privacy Workshops (SPW)*. IEEE, May 2018. DOI: 10.1109/spw.2018.00042. URL: https://doi.org/10.1109/spw.2018.00042.

[9] Timibloudi S Enamamu et al. "Smart watch based body-temperature authentication". In: *2017 International Conference on Computing Networking and Informatics (ICCNI)*. IEEE, Oct. 2017. DOI: 10.1109/iccni.2017.8123790. URL: https://doi.org/10.1109/iccni.2017.8123790.

[10] Ge Peng et al. "Continuous Authentication With Touch Behavioral Biometrics and Voice on Wearable Glasses". In: *IEEE Transactions on Human-Machine Systems* 47.3 (June 2017), pp. 404–416. DOI: 10.1109/thms.2016.2623562. URL: https://doi.org/10.1109/thms.2016.2623562.

[11] Carmen Camara et al. "Real-time electrocardiogram streams for continuous authentication". In: *Applied Soft Computing* 68 (July 2018), pp. 784–794. DOI: 10.1016/j.asoc.2017.07.032. URL: https://doi.org/10.1016/j.asoc.2017.07.032.

[12] Lorena Gonzalez-Manzano, Jose M. De Fuentes, and Arturo Ribagorda. "Leveraging User-related Internet of Things for Continuous Authentication". In: *ACM Computing Surveys* 52.3 (May 2020), pp. 1–38. DOI: 10.1145/3314023. URL: https://doi.org/10.1145/3314023.

[13] Gabriel Dahia, Leone Jesus, and Mauricio Pamplona Segundo. "Continuous authentication using biometrics: An advanced review". In: *WIREs Data Mining and Knowledge Discovery* 10.4 (Mar. 2020). DOI: 10.1002/widm.1365. URL: https://doi.org/10.1002/widm.1365.

[14] Yunji Liang et al. "Behavioral Biometrics for Continuous Authentication in the Internet-of-Things Era: An Artificial Intelligence Perspective". In: *IEEE Internet of Things Journal* 7.9 (Sept. 2020), pp. 9128–9143. DOI: 10.1109/jiot.2020.3004077. URL: https://doi.org/10.1109/jiot.2020.3004077.

[15] Sultan Almalki, Nasser Assery, and Kaushik Roy. “An Empirical Evaluation of Online Continuous Authentication and Anomaly Detection Using Mouse Clickstream Data Analysis”. In: *Applied Sciences* 11.13 (June 2021), p. 6083. DOI: `10.3390/app11136083`. URL: `https://doi.org/10.3390/app11136083`.

[16] Ines Brosso et al. “A Continuous Authentication System Based on User Behavior Analysis”. In: *2010 International Conference on Availability, Reliability and Security*. IEEE, Feb. 2010. DOI: `10.1109/ares.2010.63`. URL: `https://doi.org/10.1109/ares.2010.63`.

[17] Antonia Azzini et al. “A fuzzy approach to multimodal biometric continuous authentication”. In: *Fuzzy Optimization and Decision Making* 7.3 (June 2008), pp. 243–256. DOI: `10.1007/s10700-008-9034-1`. URL: `https://doi.org/10.1007/s10700-008-9034-1`.

[18] David Crouse et al. “Continuous authentication of mobile user: Fusion of face image and inertial Measurement Unit data”. In: *2015 International Conference on Biometrics (ICB)*. IEEE, May 2015. DOI: `10.1109/icb.2015.7139043`. URL: `https://doi.org/10.1109/icb.2015.7139043`.

[19] Andrea F. Abate, Michele Nappi, and Stefano Ricciardi. “I-Am: Implicitly Authenticate Me—Person Authentication on Mobile Devices Through Ear Shape and Arm Gesture”. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 49.3 (Mar. 2019), pp. 469–481. DOI: `10.1109/tsmc.2017.2698258`. URL: `https://doi.org/10.1109/tsmc.2017.2698258`.

[20] Attaullah Buriro, Bruno Crispo, and Yury Zhauniarovich. “Please hold on: Unobtrusive user authentication using smartphone’s built-in sensors”. In: *2017 IEEE International Conference on Identity, Security and Behavior Analysis (ISBA)*. IEEE, Feb. 2017. DOI: `10.1109/isba.2017.7947684`. URL: `https://doi.org/10.1109/isba.2017.7947684`.

[21] Yaniv Taigman et al. “DeepFace: Closing the Gap to Human-Level Performance in Face Verification”. In: *2014 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, June 2014. DOI: `10.1109/cvpr.2014.220`. URL: `https://doi.org/10.1109/cvpr.2014.220`.

[22] Hui Zhang et al. "PaddleSpeech: An Easy-to-Use All-in-One Speech Toolkit". In: *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*. Association for Computational Linguistics, 2022. DOI: `https://doi.org/10.48550/arXiv.2205.12007`.

[23] Thang Hoang, Deokjai Choi, and Thuc Dinh Nguyen. "On the Instability of Sensor Orientation in Gait Verification on Mobile Phone". In: *SECRYPT 2015 - Proceedings of the 12th International Conference on Security and Cryptography, Colmar, Alsace, France, 20-22 July, 2015.* 2015, pp. 148–159. DOI: `10.5220/0005572001480159`. URL: `https://doi.org/10.5220/0005572001480159`.

[24] *Java*. URL: `https://www.java.com`.

[25] *Python3*. URL: `https://www.python.org`.

[26] *MATLAB R2022a*. The Mathworks, Inc. Natick, Massachusetts, 2022. URL: `https://it.mathworks.com/products/matlab.html`.

[27] *Google Activity Recognition API*. URL: `https://developers.google.com/location-context/activity-recognition`.