

Received 10 November 2023, accepted 17 December 2023, date of publication 21 December 2023, date of current version 4 January 2024.

Digital Object Identifier 10.1109/ACCESS.2023.3345452

RESEARCH ARTICLE

Keystroke Verification Challenge (KVC): Biometric and Fairness Benchmark Evaluation

GIUSEPPE STRAGAPEDE¹, RUBEN VERA-RODRIGUEZ¹, RUBEN TOLOSANA¹,
AYTHAMI MORALES¹, NASER DAMER², (Member, IEEE),
JULIAN FIERREZ¹, (Member, IEEE), AND JAVIER ORTEGA-GARCIA¹, (Fellow, IEEE)

¹Biometrics and Data Pattern Analytics (BiDA) Laboratory, Universidad Autónoma de Madrid, 28049 Madrid, Spain

²Fraunhofer Institute for Computer Graphics Research IGD, 64283 Darmstadt, Germany

Corresponding author: Giuseppe Stragapede (giuseppe.stragapede@uam.es)

This work was supported in part by the European Union's Horizon 2020 Research and Innovation Program under the Marie Skłodowska-Curie under Grant 860315 (PriMa project), in part by INTER-ACTION Project under Grant PID2021-126521OB-I00 MICINN/FEDER, in part by the HumanCAIC Project under Grant TED2021-131787BI00 MICINN, and in part by the Autonomous Community of Madrid (European Laboratory for Learning and Intelligent Systems (ELLIS) Unit Madrid).

ABSTRACT Analyzing keystroke dynamics (KD) for biometric verification has several advantages: it is among the most discriminative behavioral traits; keyboards are among the most common human-computer interfaces, being the primary means for users to enter textual data; its acquisition does not require additional hardware, and its processing is relatively lightweight; and it allows for *transparently* recognizing subjects. However, the heterogeneity of experimental protocols and metrics, and the limited size of the databases adopted in the literature impede direct comparisons between different systems, thus representing an obstacle in the advancement of keystroke biometrics. To alleviate this aspect, we present a new experimental framework to benchmark KD-based biometric verification performance and fairness based on *tweet*-long sequences of variable transcript text from over 185,000 subjects, acquired through desktop and mobile keyboards, extracted from the Aalto Keystroke Databases. The framework runs on CodaLab in the form of the Keystroke Verification Challenge (KVC). Moreover, we also introduce a novel fairness metric, the Skewed Impostor Ratio (SIR), to capture *inter*- and *intra*-demographic group bias patterns in the verification scores. We demonstrate the usefulness of the proposed framework by employing two state-of-the-art keystroke verification systems, *TypeNet* and *TypeFormer*, to compare different sets of input features, achieving a less privacy-invasive system, by discarding the analysis of text content (ASCII codes of the keys pressed) in favor of extended features in the time domain. Our experiments show that this approach allows to maintain satisfactory performance.

INDEX TERMS Keystroke dynamics, behavioral biometrics, biometric verification, KVC, challenge.

I. INTRODUCTION

A. KEYSTROKE DYNAMICS FOR BIOMETRIC RECOGNITION

Keystroke Dynamics (KD) refers to the typing behavior of human subjects. It is commonly regarded as a *behavioral biometric* trait, similarly to voice [1], signature [2], [3], gait [4], [5], [6], touch gestures [7], [8], etc. In comparison

The associate editor coordinating the review of this manuscript and approving it for publication was Jad Nasreddine¹.

with its physiological counterparts such as face or fingerprint, behavioral biometrics represent a more challenging technical problem in terms of recognition performance as they are in general characterized by a higher *intra-user* variability, and lower *inter-user* variability. These challenges are magnified when dealing with real-life applications that have up to millions of subjects. Nevertheless, they offer the advantage of verifying identities *transparently*, improving security, as well as usability [9], since they spare users from having to actively carry out a specific verification procedure

(such as having their fingerprint scanned, or typing a password).

In particular, the deployment of keystroke dynamics verification systems is also economic, as there is no need for additional hardware. Potential applications span from verifying the subject identity while they write an email or they take a test in online educational platforms (free-text format) [10], to identifying malicious users across multiple accounts based on their typing style (free-text format) [11], or as an additional biometric security layer on top of a traditional knowledge-based password (fixed-text format) [12], etc. These aspects have prompted several companies to develop commercial solutions to enhance the security of users through KD [11].

A coarse classification of KD can take place according to two criteria: (i) the typology of acquisition device (keyboard): desktop or mobile. Due to differences in the pose or activity of typing subjects, more variability is commonly associated to mobile touchscreens in comparison with desktop keyboards; and (ii) regarding the text format, which can be free, fixed, or transcript. In the first case, the text typed is not the same across different samples: consequently, data are much sparser, more unstructured, and they present a higher rate of typing errors, compared to the fixed-text case, which aims to represent for instance the case of an intruder typing the password of the victim. Finally, the transcript text could be defined as a hybrid format as the subjects are asked to read, memorize, and type a text that is presented to them.

In its simplest form, keystroke dynamics are captured as discrete time instants: the time instants a key is pressed and released (for instance in Unix time format), accompanied by the code (ASCII) of the key pressed. More complex features can be extracted from these raw data. In particular, the ASCII codes are useful for learning relations between time and spatial distributions over the keyboard layout. Nevertheless, although handled in compliance with sensitive data protection regulations [13], they inevitably reveal the content of the text, putting at risk the privacy of the subjects [14], [15]. Other information such as the amount of pressure on the key or the size of the fingertip might be available depending on the specific hardware capabilities.

B. LIMITATIONS OF EXISTING EVALUATION METHODOLOGIES

In the field of keystroke biometrics, a typical obstacle for research advancement is represented by the heterogeneity of databases, experimental protocols, and metrics. In Table 1, some of the most important public keystroke dynamics databases are reported in chronological order. Although the literature on keystroke biometrics is extensive, to the best of our knowledge, except very few cases [11], [16], previous systems have mostly been only evaluated with up to several hundred subjects not representing well the

TABLE 1. Some of the most important public keystroke dynamics databases in chronological order.

Database	Scenario	No. of Subjects	Text Format	Strokes per Subject
GREYC (2009) [18]	Desktop	133	Fixed	~800
CMU (2009) [19]	Desktop	51	Fixed	~400
BiosecurID (2010) [20]	Desktop	400	Free	~200
RHU (2014) [21]	Desktop	53	Fixed	~600
Clarkson I (2014) [22]	Desktop	39	Fixed, free	~20k
SUNY (2016) [23]	Desktop	157	Transcript, free	~17k
Clarkson II (2017) [24]	Desktop	103	Free	~125k
Aalto Desktop (2018) [25]	Desktop	168k	Transcript	~750
Aalto Mobile (2019) [26]	Mobile	37k	Transcript	~750
HuMldb (2020) [27]	Mobile	600	Fixed	~20

recent challenges that massive usage applications can face. In addition, most research works are mainly focused only on desktop and fixed-text scenarios. Therefore, keystroke dynamics can still be considered a biometric modality at the early stages, especially for mobile devices. In fact, for mechanical keyboards of desktop computers, more in-depth evaluations have been conducted and commercial applications have been proposed [17]. Moreover, even if using the same databases, different systems proposed in the literature over the years have often been developed based on different subsets of users for development and evaluation, number of enrolment sessions, and metrics, hindering direct comparisons. In contrast, we propose a clearly defined experimental protocol based on same realistic use cases (Sec. VI), that can be easily adopted by researchers and practitioners of the field using the provided comparison files (Sec. II).

C. FAIRNESS CONSIDERATIONS

Moreover, in the context of decision-making, algorithms are vulnerable to biases that render their decisions “unfair” [28], [29]. Consequently, in this context, *fairness* is defined as the absence of any prejudice or favoritism toward an individual or group based on their inherent or acquired characteristics. In the last years, innumerable studies have highlighted the existence of biases in biometric systems with regard to categories such as age, gender, and ethnicity [30], leading to worse decisions that affect specific demographic groups. In addition, these aspects are also relevant from the point of view of the privacy of users, as the existence of bias due to sensitive attributes often implies that the sensitive attributes themselves might be embedded in the learned representations [31], [32]. Therefore, the risk of

leakage of some soft-biometric¹ information about the subjects should be assessed as well. In general terms, the existence of bias or privacy leakage in biometric systems presupposes specific patterns in the input data associated with different demographic groups. For instance, in face biometrics, the existence of biological differences between different genders, ages, or ethnic groups is a trivial hypothesis that does not need a formal demonstration. However, for many biometric modalities including KD, it is not straightforward to make similar assumptions. Nevertheless, for KD, several studies have evaluated the predictability of gender [35], [36], age [37], [38], both [39], [40], and even emotions [41] and mother tongue [42]. In light of this, in this article we propose an experimental framework designed to highlight potential gender and age biases in the scores, which are still mainly unexplored aspects for KD on such a large scale. Within the current work, the focus is limited to age and gender because other potential sources of bias (such as the subject mother tongue, the device used, or the degree of familiarity of the subjects with keyboards) were not reported for most subjects in the raw databases.

D. CONTRIBUTIONS

In brief, the main contributions of this article can be summarized as follows:

- We propose a novel experimental framework to benchmark KD for biometric verification, which, to the best of our knowledge, is still lacking in this field. The framework is provided in the form of the Keystroke Verification Challenge (KVC),² hosted on CodaLab.³ The CodaLab platform returns several metrics (Sec. III) that quantify the recognition performance as well as the fairness of biometric systems. To create the framework, we consider two of the largest public databases of keystroke dynamics up to date, the Aalto Desktop [25] and Mobile [26] Keystroke Databases, extracting datasets that guarantee a minimum amount of data per subject, age and gender annotations, absence of corrupted data, and that avoid too unbalanced subject distributions with respect to the considered demographic attributes.
- We illustrate the main aspects of the proposed framework by considering two recent state-of-the-art keystroke biometric systems, TypeNet [11], and TypeFormer [16], [43]. To this end, we propose a thorough analysis considering four different sets of features (Sec. VI-A) towards more privacy-preserving

biometric systems not requiring the ASCII code, which would reveal the text content, as an input feature. Our experiments show that by removing spatial information of the key location on the keyboard layout (ASCII code) in favor of additional features in the time domain, an acceptable level of performance is maintained.

- A comparative analysis of keystroke dynamics verification systems in desktop and mobile scenarios is provided (Sec. VII-A).
- We propose a new metric, the Skewed Impostor Ratio (SIR), useful to quantify how harder is for the classifier a pairwise comparison between subjects belonging to the same demographic group in relation with comparisons of subject belonging to different groups.

The remainder of the article is organized as follows: first, the resources provided within the proposed experimental framework are described (Sec. II). Then, Sec. III includes a detailed presentation of the evaluation protocol of the experimental framework and challenge, whereas Sec. IV presents the metrics adopted, including the definition of SIR, a novel metric proposed in this article. Sec. V provides an overview of the two biometric systems, TypeNet [11] and TypeFormer [16], utilized to validate the framework, followed by Sec. VI, in which the set of experiments for privacy-enhancement is illustrated. Finally, Sec. VII and Sec. VIII respectively contain the analysis of the results obtained and the article conclusive remarks.

II. RESOURCES PROVIDED

The proposed experimental framework is based on the two most complete and large-scale public databases of free-text keystroke dynamics up to date, collected by the User Interfaces⁴ group of the Aalto University (Finland). The two databases are collected respectively in a desktop⁵ [25] and mobile⁶ [26] acquisition environment, including respectively around 168,000 and 60,000 subjects, thus representing well the typical challenges related to massive application usage. Each of the acquisition sessions contains a sentence of transcript text (variable content, but not fully free-text). The data were captured through a web application in an unsupervised way under realistic scenarios. Subjects were asked to read, memorize, and type in their device English sentences that were randomly selected from a set of 1,525 sentences. Subject metadata such as age and gender are self-reported during the data acquisition.

The two databases have been processed to arrange the data in a convenient format for the analysis of KD. The raw data acquired consist of the timestamp of the instant a key is pressed, the timestamp of the instant the key is released, and the key ASCII code. After discarding some of the subject data due to insufficient acquisition sessions per subject (less than

¹Soft biometrics are physical or behavioral biometrics, or material accessories, associated with an individual, useful for recognizing an individual [33]. Examples include age, gender, ethnicity, etc. [34].

²The challenge is held within the 2023 IEEE International Conference on Big Data (IEEE Big Data), Sorrento, Italy, December 15th-18th, 2023. After such term, the challenge will be made ongoing so that the proposed framework can become a useful resource for all researchers and practitioners of the field. Website: <https://sites.google.com/view/bida-kvc/>

³<https://codalab.lisn.upsaclay.fr/competitions/14063/>

⁴<https://userinterfaces.aalto.fi/>

⁵<https://userinterfaces.aalto.fi/136Mkeystrokes/>

⁶<https://userinterfaces.aalto.fi/typing37k/>

15 per subject), the two databases as downloaded have been rearranged to form four datasets:

- Desktop Dataset:
 - 1) Development set: 115,120 subjects provided in a single.npy file that contains a Python nested dictionary (subject IDs: session IDs: data). Average session length: 48.65 ($\sigma = 18.50$) characters typed.
 - 2) Evaluation set: data from 15,000 subjects, provided in a single.npy file that contains a shallow Python dictionary (sessions IDs: data). Average session length: 48.77 ($\sigma = 18.64$) characters typed.
- Mobile Dataset:
 - 1) Development set: 40,639 subjects provided in a single.npy file that contains a Python nested dictionary (subject IDs: session IDs: data). Average session length: 48.59 ($\sigma = 21.84$) characters typed.
 - 2) Evaluation set: data from 5,000 subjects, provided in a single.npy file that contains a shallow Python dictionary (sessions IDs: data). Average session length: 47.98 ($\sigma = 20.93$) characters typed.

The proposed experimental framework follows an open-set learning protocol, in other words, the subjects in the development and evaluation sets are different⁷ in 10 subsets. Then, we computed the global EER (Sec. IV) for each of the random subsets, to provide mean and standard deviation. As an example, we report the following values for TypeFormer 5F (Sec. VI-A): $\mu = 12.949\%$, $\sigma = 0.090\%$ for desktop, $\mu = 10.164\%$, $\sigma = 0.073\%$ for mobile. Please refer to Sec. IV and VII for details about the metrics and results. A validation set is not explicitly provided, but it can be obtained from the development set according to different training approaches.

Table 2 shows the demographic distribution of the datasets provided in the KVC. The subjects have been divided into six age groups (10 - 13, 14 - 17, 18 - 26, 27 - 35, 36 - 44, 45 - 79). The evaluation sets are balanced with respect to gender. The gender and age labels of the development set are released together with the data.

The evaluation sets are separated by scenario (desktop and mobile), and they are provided in the form of two shallow Python dictionaries containing independent sessions. Such data are accompanied by the respective lists of pairwise comparisons to be carried out. Two Python script files are provided to load the data, and run the comparisons, generating a text file with the scores of each comparison, ready to be submitted to CodaLab for scoring. To push forward the state of the art and deepen the knowledge on the topic, the proposed protocol is designed for researchers working on KD as a novel tool to evaluate different approaches (pre-processing of input features, classifier architectures, learning approaches, etc.) for different goals (biometric recognition and fairness improvement)

⁷As the datasets are fixed with a single choice of subjects, the conclusions obtained from the reported results can be impacted. To verify that the score fluctuations are not significant, we randomly split the evaluation score lists (Sec.

TABLE 2. Demographic distributions of the provided datasets. The rows represent different age groups, while the columns represent genders. The evaluation sets are balanced with respect to gender.

Task 1: Desktop Dataset					
Development Set			Evaluation Set		
	Male	Female		Male	Female
10 - 13	4,336	5,420	10 - 13	1,085	1,085
14 - 17	10,993	8,336	14 - 17	1,861	1,861
18 - 26	25,752	24,315	18 - 26	1,861	1,861
27 - 35	9,607	12,281	27 - 35	1,861	1,861
36 - 44	2,143	5,331	36 - 44	536	536
45 - 79	1,182	5,424	45 - 79	296	296
Total Labelled: 115,120, Total unlabeled: 0			Total Labelled: 15,000, Total unlabeled: 0		

Task 2: Mobile Dataset					
Development Set			Evaluation Set		
	Male	Female		Male	Female
10 - 13	622	800	10 - 13	254	254
14 - 17	1,537	1,516	14 - 17	618	618
18 - 26	4,359	8,999	18 - 26	843	843
27 - 35	1,343	4,002	27 - 35	547	547
36 - 44	382	1,333	36 - 44	156	156
45 - 79	200	739	45 - 79	82	82
Total Labelled: 25,832, Total unlabeled: 14,807 ⁸			Total Labelled: 5,000, Total unlabeled: 0		

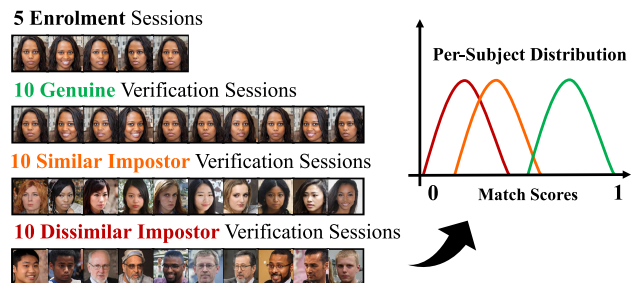


FIGURE 1. Each one of the verification sessions is compared with each of the enrolment sessions. For an easier comprehension, examples of faces showing gender and age are included instead of keystroke examples. Then, the scores generated are averaged over the enrolment session, leading to three distributions: genuine (green), similar impostor (orange), and dissimilar impostors (red).

under the same experimental conditions, considering various metrics.

III. EVALUATION DESCRIPTION

The design and the implementation of the evaluation protocol described in this section represents a significant novelty aspect proposed in the current work.

The two tasks (desktop and mobile) are structured similarly, and they are designed for a biometric verification protocol. In other words, a score between 0 and 1 related to a single comparison of two biometric samples will be produced (1: same identity, 0: different identities). It is a binary classification problem, as it is not necessary to ascertain to which identity a specific biometric sample belongs to

(identification). In this experimental framework, a biometric sample corresponds to an acquisition session.

The total number of 1 vs 1 session-level comparisons is as follows:

- Task 1 (Desktop): 2,250,000 comparisons, involving 15,000 subjects not included in the development set.
- Task 2 (Mobile): 750,000 comparisons, involving 5,000 subjects not included in the development set.

The design of the comparisons is illustrated in Fig. 1. For each subject, there are 5 enrolment sessions and 10 verification sessions, leading to 50 1vs1 comparisons, which are averaged over the 5 enrolment sessions generating 10 genuine scores per subject. In a similar manner, 20 impostor scores per subject are generated. The impostor sessions are divided into two groups: 10 *similar* impostor scores, for which the verification sessions are randomly selected from subjects belonging to the same demographic group (same gender and age); 10 *dissimilar* impostor scores, in which the verification sessions are all randomly selected from subjects of different gender and age intervals ⁹IV): ¹⁰ $\mu = 12.410\%$, $\sigma = 0.036\%$ in the desktop case, $\mu = 9.444\%$, $\sigma = 0.045\%$ in the mobile case. Please refer to Sec. IV and VII for details about the metrics and results.

Based on the described evaluation design, following [11] and [16], we consider two cases for evaluating the system:

- Global distributions: this case corresponds to dividing all scores into two groups, genuine and impostor scores, regardless of which subject they belong to. This case corresponds to a having a fixed, pre-determined threshold, implying a simpler deployment of the biometric system. In order to assess the performance of the biometric system, this choice means setting one single threshold for all comparisons to obtain a decision.
- Mean per-subject distributions: the optimal threshold is computed at subject-level, considering the 30 verification scores as described above. This choice corresponds to providing the system with more flexibility, so that it can adapt to user-specific distributions [44], [45]. In a real-life use case, this would require processing the subject's enrolment samples to establish a threshold, and it can be done as follows: acquiring various enrollment samples, from which to derive a genuine subject-specific score distribution by considering pairwise comparisons between enrollment samples; considering a pool of samples from different subjects, from which to derive an impostor subject-specific score distribution by considering pairwise comparisons with the genuine enrollment samples; computing a subject-specific threshold based on the two distributions. It is important to highlight that this does not require re-training or fine-tuning the

⁹To verify that the score fluctuations due to randomness are not significant, we repeated the choice 10 times with 10 different seeds. As an example, we report the following evaluation values for TypeFormer 5F (Sec. VI-A) in terms of global EER (Sec.

¹⁰Although unlabeled, we opted to include these subjects to maximize the size of the provided dataset.

biometric system using subject-specific data. Then, all metrics computed per-subject are averaged considering all subjects in the evaluations set to obtain the values displayed. Generally, the verification performance of the system benefits from considering a different threshold per user.

IV. METRICS ADOPTED

Within the years, several metrics have been proposed for biometric verification. The common aspect of all metrics is that they are based on the (normalized) scores that are typically generated by pairwise comparisons of biometric data. However, a comparison between systems is often a difficult operation if they are evaluated according to different metrics. Moreover, the attention of the scientific community has recently shifted towards the evaluation of the *fairness* of systems [46]. Consequently, based on the scores, we also provide an initial assessment of this important aspect which, to the best of our knowledge, is still an unexplored aspect of KD. An overview of all metrics considered is provided below. The scores will be computed considering global and mean per-subject distributions. All the presented metrics are returned by the KVC CodaLab scoring program, easily allowing experimental analyses of multiple aspects, such as different sets of input features for privacy-enhancement (Sec. VII). To the best of our knowledge, these scenarios have not been proposed in previously existing literature.

A. VERIFICATION METRICS

In biometrics, a false match (FM) is defined as a comparison decision of a match for a biometric probe and a biometric reference that are from different biometric capture subjects, while a false non-match (FNM) is defined as comparison decision of non-match for a biometric probe and a biometric reference that are from the same biometric capture subject and of the same biometric characteristic. The rate respectively associated with FMR (FNMR) corresponds to proportion of the completed biometric non-mated (mated) comparison trials that result in a false match (non-match) [47].

1) EQUAL ERROR RATE (EER)

The EER describes the point in which the FMR and FNMR curves intersect. The two rates typically have opposite trends with respect to the threshold setting (in the case of genuine scores closer to 1, and impostor scores closer to 0, as the threshold of a biometric system increases, the FMR will drop and the FNMR curve will rise). On the DET curve (Sec. IV-B), which is the plot of FNMR against FMR, at various threshold settings, it corresponds to the point where $y = x$.

2) FALSE NON-MATCH RATE AT X% FALSE MATCH RATE (FNMR @ X% FMR)

We consider $X = 1\%$, 10% . This also corresponds to a point on the DET curve. This metric expresses a trade-off between security and usability [9]. In fact, while from the

point of view of security the priority is avoiding intrusions, denying the access to the genuine subject a large number of times would generate frustration and highly impacts the usability of the system. In this case, the threshold is set to $X = 1\%$, 10% of FMR (rejection of 99%, 90% of impostor attempts, respectively), aiming to minimize the FNMR.

3) AREA UNDER THE RECEIVER OPERATING CHARACTERISTIC (ROC) CURVE (AUC)

The ROC curve (Sec. IV-B) is the plot of the TMR (True Match Rate) against FMR, at various threshold settings. A true match corresponds to the case of a genuine subject recognized as such. By definition, the TMR and the FNMR sum to 1. A perfect classifier has an Area Under the ROC Curve (AUC) of 1.

4) ACCURACY

The accuracy is computed as the fraction of correctly classified attempts at a given discrimination threshold τ , corresponding, in the current work, to the EER threshold.

5) RANK-N

This metric concerns the identification of subjects (i.e., 1 to many comparisons), therefore assessing a different scenario from the previous metrics, which refer to the case of verification (i.e., binary classification). Starting from the comparison of biometric enrolment samples with N biometric samples including a genuine one, it represents the rate to which the genuine scores fall within the best n matches. In the proposed framework, this metric is computed by considering separately each one of the genuine verification sessions against all 20 impostor sessions available for each subject. The returned rank- n values are averaged over the 10 genuine verification sessions ($n = 1$).

B. CURVES

1) SCORE HISTOGRAMS

They are computed considering the global genuine and impostor distributions. It is necessary to have a clear separation between the two, with the genuine distribution shifted toward 1, and the impostor one toward 0. A small overlap of the tails corresponds to a better performance of the system.

2) DETECTION ERROR TRADE-OFF (DET) CURVE

It is the plot of FMR against FNMR, at various threshold settings, typically on a non-linear scale. As the threshold decreases, the amount of false matches (impostor subjects classified as genuine) increases, and the number of false non-matches decreases (genuine subjects classified as impostor). The closer the DET curve to the bottom left corner, the better the biometric system will be.

3) ROC CURVE

It is the plot of the TMR against FMR, at various threshold settings.

C. FAIRNESS METRICS

1) STANDARD DEVIATION (STD) OF EER BY DEMOGRAPHIC GROUP

It considers the demographic differential assessment by calculating the standard deviation in accuracy performance between all demographic groups at a given discrimination threshold τ (in this work corresponding to the global EER threshold). The STD expresses a measure of how dispersed the values is in relation to the mean. The optimal STD value is 0%.

2) SKEWED ERROR RATIO (SER) OF EER BY DEMOGRAPHIC GROUP

Skewness is a measure of the asymmetry of a distribution. Similarly to the STD, SER is computed across demographic subsets as the ratio between the greatest and smallest error scores. It mainly represents the difference between the sensitive attribute with the best and worst performance. The larger the value, the greater the difference in the algorithm's discrimination towards a certain attribute. The optimal SER value is 1.

3) FAIRNESS DISCREPANCY RATE (FDR)

It was proposed in [48]. It considers the FMR and FNMR trade-off in the demographic differential assessment by calculating the max difference in TMR and TNMR performance between any two demographic groups d_i and d_j and a given discrimination threshold τ . In the current work, τ corresponds to the point of $FMR = 1\%$. Those differences are then weighted by parameters α and $\beta = 1 - \alpha$, which represent the level of concern applied to differences in FMR and FNMR respectively. It is advantageous due to its formulation, that encompasses two different metrics, and its flexibility given by the parameters α and β . It can be plotted as a function of every possible threshold value. Consequently, it is necessary to have access to the ground truth. The FDR ranges from 0 (fair) to 1 (unfair).

4) INEQUITY RATE (IR)

It was proposed in [49]. It is computed considering the ratio differences between minimum and maximum FMR and FNMR per demographic groups d_i and d_j for τ corresponding to the point of $FMR = 1\%$. Although spanning different ranges from, its characteristics are similar to those of FDR. The IR ranges from 0 (fair) to infinite (unfair).

5) GINI AGGREGATION RATE FOR BIOMETRIC EQUITABILITY (GARBE)

It was proposed in [50] to overcome the limitations of FDR, and IR. In fact, the former does not scale the values of FMR and FNMR to the same order of magnitude, whereas the latter

has no theoretical upper bound and may have a denominator equal to zero. GARBE is inspired in the mathematics of the Gini coefficient, computed for $x = \{\text{FMR}, \text{FNMR}\}$ as:

$$G_x(\tau) = \left(\frac{n}{n-1} \right) \left(\frac{\sum_{i=1}^n \sum_{j=1}^n |r_i - r_j|}{2n^2 \bar{r}} \right)$$

where $r_i, r_j \in \{\text{FMR}_{d_i} @ \tau | d_i \in \mathbb{D}\}$ for $r = \text{FMR}$, $r_i, r_j \in \{\text{FNMR}_{d_i} @ \tau | d_i \in \mathbb{D}\}$ for $r = \text{FNMR}$, τ is the decision threshold ($\text{FMR} = 1\%$), d_i refers to a demographic group from the set \mathbb{D} of demographic groups, and n is the number of demographic groups. These Gini coefficients are combined as follows:

$$\text{GARBE}(t_z) = \alpha G_{\text{FMR}}(t_z) + (1 - \alpha) G_{\text{FNMR}}(t_z)$$

where α represents the level of concern applied to differences in FMR and FNMR respectively. The GARBE ranges from 0 (fair) to 1 (unfair).

6) SKEWED IMPOSTOR RATIO (SIR)

This is a novel metric proposed in this article. Normalized impostor scores are grouped according to a specific attribute (age or gender). For instance, considering age, it is possible to group the comparisons as follows: ‘10-13 vs 10-13’, ‘10-13 vs 14-17’, ‘10-13 vs 18-26’, and so on, considering all combinations. For each score group combination, the average value is taken. Then, all values can be arranged in a matrix, which is symmetric. The elements on the main diagonal represent comparisons between the different subjects belonging to the the same age or gender group (‘10-13 vs 10-13’, ‘14-17 vs 14-17’, etc.), while all other elements are obtained from the remaining cross-group comparisons. The ratio between the mean value of the elements in the main diagonal, and the remaining non-duplicated elements, is finally computed as a percentage. Such value expresses how harder is a comparison between different subjects belonging to the same demographic group in comparison to subjects belonging to different ones, quantifying to which extent demographic information is retained in the scores. It can be formulated as follows:

$$\text{SIR} = 100 \left(\frac{\mu(s_{ii})}{\mu(s_{ij, i \neq j})} - 1 \right), i, j = 1, \dots, n$$

where n is the number of demographic groups ($n = 6$ for age, $n = 2$ for gender), and s is the average score for a specific type of comparisons. The optimal SIR value is 0 (%), in case of no skew between impostor scores regardless of their demographic group. In comparison with the previous metrics, the advantages are as follows:

- It is not necessary to select a threshold value.
- It focuses on both *intra-group* and *inter-group* relations, highlighting the differences in the two cases. If the differences between the two cases are not significant nor consistent, then the system is bias-free.
- It is not necessary to have access to the system, which can be treated as a *black box*. It is sufficient

to run a significant number of appropriately distributed comparisons.

- By considering the entire matrix as described above, it is possible to focus on comparisons between specific groups, gaining some precious insights about the system and the similarities between demographic groups.

V. BIOMETRIC VERIFICATION SYSTEMS

Throughout the proposed framework, we evaluate two recent state-of-the-art deep-learning models:

- TypeNet (2021) [11]: a Long-Short Term Memory (LSTM) Recurrent Neural Network (RNN), trained with triplet loss. In this case, we consider input sequences of 150 characters typed. TypeNet is implemented in Tensorflow [51].
- TypeFormer (2023) [16], [43]: a novel transformer architecture consisting in a temporal and a channel module enclosing two LSTM RNN layers, a Gaussian Range Encoding (GRE), a multi-head self-attention mechanism, and a block-recurrent transformer structure. TypeFormer is also trained with triplet loss. In this case, we consider input sequences of 50 characters typed. TypeFormer is implemented in PyTorch [52].

Both approaches utilize Distance Metric Learning (DML) [53]. The fundamental concept of DML involves training a model that transforms input data into a new feature space, enabling straightforward distances to be used for analyzing and leveraging the “semantic” arrangement of the input space [54]. A DML approach aims to establish a neighborhood structure in the feature space by considering the relationship between intra-class (distances among samples from the same class) and inter-class (distances among samples from different classes) distances. In an ideal feature space, samples from the same class will remain in close proximity, while samples from different classes will be distinctly separated. Following this idea, the input sequences obtained from all sessions are transformed into feature embeddings, that are expected to have lower Euclidean distances if belonging to the same subject, higher otherwise. In the test stage, the distances obtained from the comparisons of feature embeddings corresponding to each of the test sessions are normalized, and then they are subtracted from 1 in order to transform them into similarity scores.

VI. EXPERIMENTAL PROTOCOL

A. EVALUATION OF PRIVACY-ENHANCING INPUT FEATURES

The two biometric systems considered, TypeNet and TypeFormer, take as input the same set of features extracted from the raw data (Unix timestamps of the actions of pressing and releasing a key), which include:

- Hold Time (HT): time interval between the release and press instants of a given key, expressed in seconds.
- Inter-Press Time (IPT): time interval between two consecutive press actions, expressed in seconds.

(ii) Inter-Release Time (IRT): time interval between two consecutive release actions, expressed in seconds.

(iv) Inter-Key Time (IKT): time interval between a release and the following press action, expressed in seconds.

(v) ASCII code (ASCII), normalized by dividing it by 255. Such features are graphically represented in black in Fig. 2. Input features (i)-(iv) are useful to capture the typing behavior of the user in the time-domain, whereas the ASCII code (v) describes the spatial relations due to the location of the key pressed on the keyboard layout. However, although handled in compliance with sensitive data protection regulations [13], the acquisition and processing of the ASCII codes inevitably reveals the content of the text, putting at risk the privacy of the users. Consequently, in this experiment, we strive to remove the ASCII code information in order to make the system content-agnostic, and consequently more privacy preserving. The sets of experiments run can be summarized as follows:

1. First, to evaluate and compare TypeNet and TypeFormer, we consider their original set of features (the experiment is named $5F$, where “ F ” stands for “features”), marked in black in Fig. 2.
2. Then to quantify the importance of the ASCII code information, we remove the ASCII code information from the original set of features (experiment $4F$).
3. We consider an extended set of time-domain features. Several studies have in fact shown the usefulness of considering groups of keys typed such as digraphs, trigraphs, and n -graphs [55]. By considering not only adjacent keys, but groups of three keys (Fig. 2, in red and blue), we obtain a set of 10 features (experiment $10F$):

$$[\text{HT}, \text{IPT}, \text{IRT}, \text{IKT}, \text{IPT2}, \text{IRT2}, \text{IKT2}, \text{IPT3}, \text{IRT3}, \text{IKT3}]$$

4. We consider the extended set of time-domain features, together with the ASCII code (experiment $11F$).
- In each case, the deep learning models are trained from scratch.

B. MODEL TRAINING

The training of both models takes place on the KVC development set considering identical settings to those described in their respective papers. The only differences are related to the division into training and validation sets. For TypeNet, we consider a subset of 400 subjects to validate the model at the end of each training epoch in terms of average EER per subject. This choice is justified by the experimental protocol followed in [11]. For TypeFormer, we consider an 80%-20% train-validation division of the KVC development set, and we adopt the global EER as validation metric. According to these validation metrics, the best-performing epoch model is saved in each case.

VII. EXPERIMENTAL RESULTS

A. BIOMETRIC VERIFICATION

The results of the experiments are reported in Table 3. The table is divided into two parts, each one corresponding to

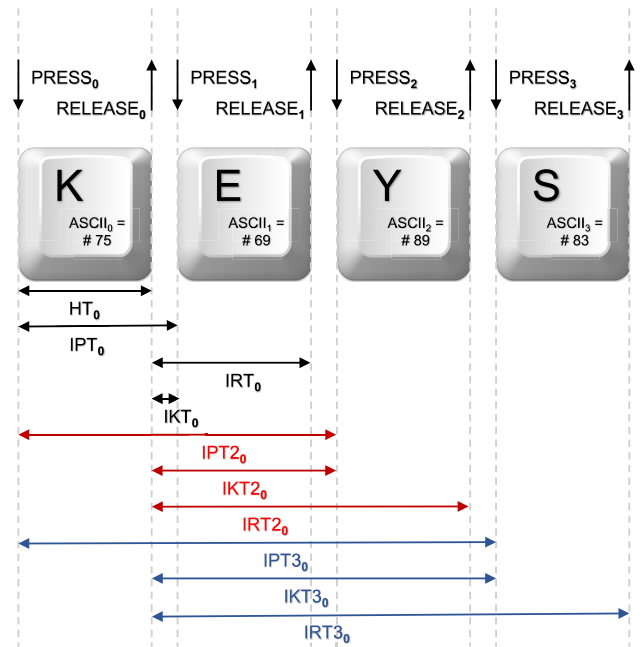


FIGURE 2. A diagram representing the initial feature extraction process for the time instant $t = 0$. HT = Hold Time; IPT = Inter-Press Time; IRT = Inter-Release Time; IKT = Inter-Key Time; IPT2 = Inter-Press Time with second following key; IRT2 = Inter-Release Time with second following key; IKT2 = Inter-Key Time with second following key; IPT3 = Inter-Press Time with third following key; IRT3 = Inter-Release Time with third following key; IKT3 = Inter-Key Time with third following key.

one scenario: desktop and mobile. Each half can be further divided into the two cases considered: results obtained in the global genuine and impostor distributions (see Sec. III), and results considering the mean values obtained for per-subject genuine and impostor distributions. Each row shows a different system, TypeNet or TypeFormer, trained on a different set of input features, according to the 4 experiments described in Sec. VI-A, while the different metrics are reported along the columns.

By observing the overall trends in the table, it is possible to notice that in the desktop case higher verification results can be achieved, possibly due to a more constrained acquisition scenario, as, in contrast to mobile devices, subjects are more likely to be sitting down and in a still position while typing on a desktop keyboard. In fact, as an example, considering both TypeNet and TypeFormer with all possible sets of features (8 experiments in total), the mean value of all EERs obtained from the global distributions in the desktop scenario is 10.53%, whereas the corresponding value obtained in the mobile case is 13.11%. The trend is consistent if we analyze the other metrics, such as the FNMR @1% FMR (60.47% vs 72.95%) or AUC (95.73% vs 93.99%). Furthermore, the desktop case results to perform better also in the case of mean per-subject distributions. In this case, we obtain a mean EER of 5.67% in the desktop case vs 7.61% in the mobile case. Similarly, the mean AUC is 97.70% vs 96.59%, and the mean rank-1

TABLE 3. Complete comparison of the presented keystroke biometric verification systems.

Desktop					
Global Distributions					
Experiment	EER (%)↓	FNMR @1% FMR (%)↓	FNMR @10% FMR (%)↓	AUC (%)↑	Accuracy (%)↑
TypeNet, 5F	6.76	39.57	3.45	98.08	93.24
TypeNet, 4F	8.97	52.17	7.56	96.92	91.03
TypeNet, 10F	8.95	51.57	7.51	96.91	91.05
TypeNet, 11F	6.88	40.0	3.64	98.02	93.12
Average	7.89	45.83	5.54	97.48	92.11
TypeFormer, 5F	12.95	75.11	19.03	94.1	87.05
TypeFormer, 4F	14.05	76.28	22.22	93.37	85.95
TypeFormer, 10F	12.75	73.51	18.19	94.31	87.25
TypeFormer, 11F	12.92	75.51	18.93	94.15	87.08
Average	13.17	75.1	19.59	93.98	86.83
Mean Per-Subject Distributions					
Experiment	EER (%)↓	AUC (%)↑	Accuracy (%)↑	Rank-1 (%)↑	
TypeNet, 5F	2.71	99.26	95.31	89.81	
TypeNet, 4F	4.13	98.67	94.29	83.46	
TypeNet, 10F	4.16	98.66	94.27	83.58	
TypeNet, 11F	2.76	99.24	95.29	89.52	
Average	3.44	98.96	94.79	86.59	
TypeFormer, 5F	7.76	96.51	91.13	65.52	
TypeFormer, 4F	8.62	95.99	90.4	62.7	
TypeFormer, 10F	7.4	96.77	91.47	67.72	
TypeFormer, 11F	7.78	96.48	91.13	65.07	
Average	7.89	96.44	91.03	65.25	
Mobile					
Global Distributions					
EER (%)↓	FNMR @1% FMR (%)↓	FNMR @10% FMR (%)↓	AUC (%)↑	Accuracy (%)↑	
TypeNet, 5F	13.95	70.05	22.22	93.8	86.05
TypeNet, 4F	16.36	76.64	28.52	91.41	83.64
TypeNet, 10F	14.51	77.07	23.79	92.71	85.49
TypeNet, 11F	14.8	68.16	23.48	93.71	85.2
Average	14.9	72.98	24.5	92.91	85.1
TypeFormer, 5F	10.17	70.57	10.46	95.86	89.83
TypeFormer, 4F	13.7	78.5	21.48	93.4	86.3
TypeFormer, 10F	9.45	67.67	8.53	96.22	90.55
TypeFormer, 11F	11.95	74.96	16.16	94.81	88.05
Average	11.32	72.92	14.16	95.07	88.68
Mean Per-Subject Distributions					
Experiment	EER (%)↓	AUC (%)↑	Accuracy (%)↑	Rank-1 (%)↑	
TypeNet, 5F	7.99	96.43	90.89	68.5	
TypeNet, 4F	9.64	95.46	89.5	65.13	
TypeNet, 10F	8.3	96.19	90.64	66.15	
TypeNet, 11F	7.99	96.54	91.03	68.49	
Average	8.48	96.16	90.51	67.07	
TypeFormer, 5F	5.92	97.45	92.64	71.3	
TypeFormer, 4F	8.42	96.09	90.61	63.1	
TypeFormer, 10F	5.25	97.89	93.28	75.92	
TypeFormer, 11F	7.41	96.66	91.42	66.04	
Average	6.75	97.02	91.99	69.09	

is 75.92% vs 68.08%, respectively for desktop and mobile devices.

Focusing on the four experiments involving different sets of input features, it is possible to draw some interesting conclusions. The best performing system for the desktop scenario is TypeNet, which is affected by the lack of the spatial information given by the ASCII code, and the extended set of features is not able to compensate it (6.76% EER for TypeNet 5F vs 8.95% EER for TypeNet 10F for global distributions). Nevertheless, despite the performance decrease, the system still shows competitive performance against the Transformer-based TypeFormer 5F

or TypeFormer 10F (respectively 12.95% EER and 12.75% EER for global distributions). Moreover, these tendencies are consistent considering mean per-subject distributions.

The opposite outcome can be observed in the case of TypeFormer, which is the best performing system in the mobile case: the model achieved with the extended set of features (10F) proves to be the best performing system in the mobile scenario (9.45% EER vs second best, 5F, of 10.17% EER for global distributions). Consequently, by introducing temporally-deepened input features, not only it has been possible to limit the performance decrease towards a more privacy-preserving verification system, but




FIGURE 3. The graphs show a comparison between TypeNet 5F vs TypeNet 10F in the desktop case. (a) shows the score histograms, (b) shows the DET curve, (c) shows the ROC curve.




FIGURE 4. The graphs show a comparison between TypeFormer 5F vs TypeFormer 10F in the mobile case. (a) shows the score histograms, (b) shows the DET curve, (c) shows the ROC curve.

the verification performance is even significantly improved. In this case, it must be specified that the number of heads in the attention mechanism must be a multiple of the number of input features, consequently we considered (5 for 5F and 10F, 4 for 4F, and 1 for 11F due to memory constraints). In the case of TypeFormer, it is also interesting to point out that the second best performing model corresponds to the experiment 5F, with the initial set of features.

In addition, it is noticeable that subject-specific distributions lead to better verification performance in all cases. In fact, the system benefits from gaining more flexibility by setting a different threshold per subject. It is worth to highlight that this is really the case for security systems based on KD: in a real-life use case, once the system is deployed and the subject identity is verified in some other way, it is not hard to acquire multiple samples per subject. All these subject-specific data could be used as enrolment data, building a complete behavioral profile of the subject and leading to even better performance, without necessarily having to train or fine-tune the system. This would in fact be a further possible step, that leaves additional margin of improvement [56].

Another interesting observation is that it is very distinct how in the desktop case TypeNet performs better, while TypeFormer shows superiority in the mobile case. TypeNet is based on a two-layer LSTM RNN, while TypeFormer is based on a Transformer architecture, composed of several modules, and more parameters. As an example, in the desktop case the average TypeNet performance in terms of EER, FNMR @1% FNMR, and AUC, over the four input feature experiments, is respectively 7.89%, 45.83%, 97.48%, for global distributions and 3.44% (EER), 98.96% (AUC), 86.59% (rank-1), for mean-per subject distributions. In each case, the corresponding values achieved by TypeFormer are 13.17%, 75.10%, 93.98%, (global distributions), and 7.89%, 96.44%, 65.25% (mean per-subject distributions), showing a significant gap. These trends are clearly opposite in the mobile case, where TypeFormer shows 11.32%, 72.92%, 95.07% in the case of global distributions and 6.75%, 97.02%, 69.09% for the mean per-subject distributions. In the corresponding experiments, TypeNet achieves 14.91%, 72.98%, 92.91%, for global distributions and 8.48%, 96.16%, 67.07% for the mean per-subject distributions. These trends show that LSTM RNNs seem to model well the desktop environment, while the higher variability of mobile devices

is better modelled by a Transformer, which, however, does not reach the same level of performance in absolute terms.

In addition, in the case of the global distributions, also the FNMR @10% FMR is reported. This represents a more relaxed approach, as the threshold selected is less stringent (90% of the impostors are rejected against 99% of FNMR @1% FMR). According to both these metrics, TypeNet in the desktop case achieves significantly better results in comparison with all the other configurations, and its gap with TypeFormer is much greater than for the mobile case, where the two system performance is closer. These metrics are not computed for the case of mean per-subject distributions as there are not enough scores per subject for a sufficiently fine threshold resolution. They are substituted by the Rank-1 (not determinable for global distributions as subject-dependent), which also shows a more regular behavior of TypeNet in the desktop case in comparison with all other configurations.

Fig. 3 and Fig. 4 show the curves described in Sec. IV-B. In particular, it is possible to carry out a direct comparison of TypeNet in the desktop case and TypeFormer in the mobile case considering two of the four input feature sets presented above (5F vs 10F). From left to right, the histograms of the genuine and impostor distributions are reported in Fig. 3 and 4 (a). It is possible to see that in both rows the genuine distributions are more separated for the case of 5F, while for the mobile case (TypeFormer), the 10F setup is able to create a better separation of impostors, leading to a lower EER value. The threshold corresponding to the EER value is marked by the black line. From these macroscopic trends, the difference between the two impostor scenarios (“similar impostors” and “dissimilar impostors”, for comparison between subjects of the same and different demographic groups, respectively) is not very pronounced. In all graphs, the threshold values are reported in black. Then, Fig. 3 and 4 (b) represent the DET curves (the threshold corresponding to 1% and 10% of FMR are respectively marked by the dashed red lines), while Fig. 3 and 4 (c) report the ROC curves, showing similar trends from different perspectives.

B. FAIRNESS EVALUATION

Table 4 shows the performance in terms of accuracy based on the global EER threshold (%) considering different demographic groups. Age groups are placed along the rows, while genders along the columns. As an example, from all experiments we take into consideration TypeNet in the desktop case, and TypeFormer in the mobile case, based on the same set of features (“5F”). In both cases, it is possible to see that males achieved higher values (93.13% against 92.49% of global EER for TypeNet, and 89.80% against 88.55% for TypeFormer, respectively for males and females). It is interesting to notice that, although the STD and SER values are quite smaller in the desktop case, the difference of EER between error rates across genders is still

TABLE 4. Results in terms of global accuracy for the final evaluation datasets (desktop and mobile), evaluated for the different demographic groups (gender and age). STD refers to the standard deviation, whereas SER refers to the Skewed Error Rate (Sec. IV). Both metrics are computed from all the elements of each table.

TypeNet, 5F, Desktop, Global Accuracy (%)			
	Male	Female	Average
10 - 13	93.75	92.88	93.31
14 - 17	93.1	92.84	92.97
18 - 26	93.60	93.16	93.38
27 - 35	92.85	92.37	92.61
36 - 44	92.44	92.07	92.25
45 - 79	93.02	91.65	92.33
Average	93.13	92.49	-
STD: 0.603%, SER: 1.023			
TypeFormer, 5F, Mobile, Global Accuracy (%)			
	Male	Female	Average
10 - 13	88.63	87.33	87.98
14 - 17	89.04	86.52	87.78
18 - 26	90.69	88.39	89.54
27 - 35	89.37	90.05	89.71
36 - 44	90.1	90.57	90.34
45 - 79	90.94	88.43	89.68
Average	89.8	88.55	-
STD: 1.38%, SER: 1.051			

significant. Formulating an hypothesis that would explain this trend is not immediate, and out of the scope of the current work. Furthermore, by analyzing the behavior of the systems considering different age groups, it is possible to notice that for younger subjects, TypeFormer performs worse than for older ones, while this trend is not as evident for TypeNet. Such discrepancy could be due to cultural differences related to the degree of easiness and comfort in interacting with mobile devices across different generations.

Table 5 shows all results of the fairness assessment provided throughout the proposed framework. Along the rows, the table is divided into two parts: the upper part presents the results in the desktop scenario, while the lower part is focused on the mobile scenario. Each half is further divided into two sections, each reporting results of one of the two models considered, TypeNet and TypeFormer, considering the four experimental configurations presented in Sec. VI-A. Concerning the different metrics, it is necessary to point out that all metrics except SIR are calculated considering 12 demographic groups, due to 6 age groups and 2 genders (see Sec. II). In contrast, SIR_a is computed considering scores divided by age only (square matrix of dimension 6), while SIR_g considering gender only (square matrix of dimension 2), to keep the assessments of each of the two attributes independent (Sec. IV).

By observing Table 5, it is possible to notice that there is no clear superiority of one model as in the case of the verification performance (TypeNet 5F for desktop devices, TypeFormer 10F for mobile devices). By carrying out an overall comparison of the desktop and mobile scenarios, it is

TABLE 5. Comparison of the presented keystroke biometric verification systems from the point of view of fairness metrics.

Desktop							
Experiment	STD (%)↓	SER↓	FDR↑	IR↓	GARBE↓	SIR _a (%)↓	SIR _g (%)↓
TypeNet, 5F	0.603	1.023	95.234	1.39	0.056	2.919	2.229
TypeNet, 4F	0.72	1.03	93.267	1.43	0.06	2.625	2.121
TypeNet, 10F	0.763	1.033	94.849	1.434	0.062	2.623	2.024
TypeNet, 11F	0.694	1.028	94.045	1.528	0.073	3.046	2.252
TypeFormer, 5F	0.817	1.034	97.562	1.296	0.05	3.137	2.649
TypeFormer, 4F	0.863	1.038	97.649	1.185	0.032	2.97	2.452
TypeFormer, 10F	0.847	1.036	96.71	1.202	0.037	2.973	2.493
TypeFormer, 11F	0.851	1.034	97.882	1.405	0.059	3.161	2.521
Mobile							
Experiment	STD (%)↓	SER↓	FDR↑	IR↓	GARBE↓	SIR _a (%)↓	SIR _g (%)↓
TypeNet, 5F	2.234	1.092	95.727	2.089	0.099	6.714	8.305
TypeNet, 4F	1.95	1.077	94.878	3.413	0.175	2.883	3.518
TypeNet, 10F	2.037	1.078	94.948	2.451	0.17	4.205	5.15
TypeNet, 11F	1.679	1.074	95.202	2.36	0.1	7.076	8.722
TypeFormer, 5F	1.38	1.051	95.015	2.113	0.138	6.901	7.569
TypeFormer, 4F	1.764	1.054	96.042	2.521	0.142	6.428	6.262
TypeFormer, 10F	1.551	1.057	94.836	2.338	0.162	6.44	6.097
TypeFormer, 11F	1.721	1.075	93.806	2.247	0.119	7.601	7.4

possible to observe that demographic differentials related to age and gender tend to be higher in the mobile case. In fact, considering the mean values for each one of the metrics computed over all desktop against all mobile experiments, we obtain STD 0.77% vs 1.79%, SER 1.032 vs 1.07, FDR 95.90 vs 95.06, IR 1.36 vs 2.44, GARBE 0.06 vs 0.14, SIR_a 2.80% vs 6.03%, SIR_g 2.34% vs 6.63%. It is clear that in the desktop scenario there is an improvement in fairness according to each single metric.

If we consider a comparison of the two architectures in each scenario, we can see that in the desktop case the average performance achieved by TypeNet in terms of SIR_a is 2.80% and SIR_g is 2.16%, while the scores of TypeFormer are respectively 3.06% and 2.53%, being more biased as well as less effective. A similar trend is reported for the mobile case, with TypeNet (although with lower verification performance) achieving SIR_a=5.22% and SIR_g=6.42%, against SIR_a=6.84% and SIR_g=6.83% achieved by TypeFormer.

Finally, having a look at trends due to different sets of input features within the same scenario and same model configuration, we can observe that the results are more irregular and they do not show a clear overall trend as in the previous cases. Nevertheless, considering for instance the SIR values, we can observe that generally the ASCII code is associated with higher bias in almost all experiments and scenarios, showing that comparisons among the same demographic groups are slightly harder than among different groups.

VIII. CONCLUSION AND FUTURE WORK

In this article an open experimental framework to benchmark keystroke dynamics for biometric verification is provided to the research community to alleviate the heterogeneity of the experimental protocols, metrics, and the limited size of the databases adopted in the literature. The framework is provided in the form of the Keystroke Verification

Challenge (KVC),¹¹ hosted on CodaLab,¹² and held within the 2023 IEEE International Conference on Big Data,¹³ Sorrento, Italy, December 15th-18th, 2023. After such term, the challenge will be made ongoing. The experimental framework is based on the two most complete and large-scale public databases of free-text keystroke dynamics up to date, collected respectively in a desktop and mobile acquisition environment, including keystroke data from more than 185,000 subjects overall. The proposed framework not only allows a complete assessment of the verification performance, but it also returns several metrics related to biometric fairness and bias based on the comparison scores, to gain new insights about KD. To this end, a novel metric, the Skewed Impostor Ratio (SIR), is proposed, designed to highlight *inter-* and *intra-*demographic group patterns present in the final comparison scores.

Finally, we make a first use of the proposed framework by employing two recent state-of-the-art keystroke biometric systems, TypeNet [11] and TypeFormer [16]. Besides a direct comparison of the two, that shows the superiority of the former in the desktop scenario, and of the latter in the mobile one, we consider different sets of input features towards a more privacy-preserving keystroke verification system. The proposed solution is based on discarding the ASCII code, which reveals the text content, in favor of extended features in the time domain. We analyze four different experimental configurations, that focus on deepening the temporal information provided to the models in order to compensate the removal of spatial information due to the ASCII code, utilized to learn the relation in between the location of the specific key in the keyboard layout, and the correspondent time dynamics. Our experimental results show that such approach allows to maintain satisfactory

¹¹<https://sites.google.com/view/bida-kvc/>

¹²<https://codalab.lisn.upsaclay.fr/competitions/14063/>

¹³<http://bigdataieee.org/BigData2023/>

performance in the desktop scenario, and even improved for mobile devices.

Concerning future work, the next directions of research will go toward the optimisation of the model architectures to improve the recognition performance and to reduce bias. More sophisticated training approaches will also be investigated, i.e. loss functions based on the selections of hard comparisons and adaptive margins [53], [57], specifically designed for the case of behavioral biometrics such as KD. Moreover, approaches based on the generation of synthetic subject-specific data will be considered to assess the suitability of such techniques to the problem of behavioral biometrics-based verification. To this end, KVC represents a dedicated and unified test bench for the entire biometric research community. By doing so, we aim to foster the design of innovative solutions that achieve improved performance in comparison with existing ones, that are benchmarked here.

Finally, the results of our contributed benchmark KVC in terms of demographic attribute assessment also enable further large-scale studies focused on examining the differences in subjects' typing behavior due to biological, cultural, or linguistic factors. In this sense, new findings might be of great interest for several branches of the Human-Computer Interaction (HCI) community, e.g. privacy protection [14], security of minors online [58], user experience improvement [59], etc.

REFERENCES

- [1] X. Chen, Z. Li, S. Setlur, and W. Xu, "Exploring racial and gender disparities in voice biometrics," *Sci. Rep.*, vol. 12, no. 1, p. 3723, Mar. 2022.
- [2] R. Tolosana et al., "SVC-onGoing: Signature verification competition," *Pattern Recognit.*, vol. 127, p. 108609, 2022.
- [3] R. Tolosana, R. Vera-Rodriguez, J. Fierrez, and J. Ortega-Garcia, "DeepSign: Deep on-line signature verification," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 3, no. 2, pp. 229–239, Apr. 2021.
- [4] O. Costilla-Reyes, R. Vera-Rodriguez, P. Scully, and K. B. Ozanyan, "Analysis of spatio-temporal representations for robust footstep recognition with deep residual neural networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 2, pp. 285–296, Feb. 2019.
- [5] P. Delgado-Santos, R. Tolosana, R. Guest, R. Vera-Rodriguez, and J. Fierrez, "M-GaitFormer: Mobile biometric gait verification using transformers," *Eng. Appl. Artif. Intell.*, vol. 125, Oct. 2023, Art. no. 106682.
- [6] P. Delgado-Santos, R. Tolosana, R. Guest, F. Deravi, and R. Vera-Rodriguez, "Exploring transformers for behavioural biometrics: A case study in gait recognition," *Pattern Recognit.*, vol. 143, Nov. 2023, Art. no. 109798.
- [7] G. Stragapede, R. Vera-Rodriguez, R. Tolosana, A. Morales, J. Fierrez, J. Ortega-Garcia, S. Rasnayaka, S. Seneviratne, V. Dissanayake, J. Liebers, A. Islam, S. B. Belhaouari, S. Ahmad, and S. Jabin, "IJCB 2022 mobile behavioral biometrics competition (MobileB2C)," in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Oct. 2022, pp. 1–7.
- [8] G. Stragapede, R. Vera-Rodriguez, R. Tolosana, and A. Morales, "BehavePassDB: Public database for mobile behavioral biometrics and benchmark evaluation," *Pattern Recognit.*, vol. 134, Feb. 2023, Art. no. 109089.
- [9] *Ergonomic Requirements for Office Work With Visual Display Terminals (VDTs), 1998. Part 11: Guidance on Usability*, Standard ISO 9241-11, International Organization for Standardization, Geneva, Switzerland, 1998.
- [10] J. Hernandez-Ortega, R. Daza, A. Morales, J. Fierrez, and J. Ortega-Garcia, "edBB: Biometrics and behavior for assessing remote education," 2019, *arXiv:1912.04786*.
- [11] A. Acien, A. Morales, J. V. Monaco, R. Vera-Rodriguez, and J. Fierrez, "TypeNet: Deep learning keystroke biometrics," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 4, no. 1, pp. 57–70, Jan. 2022.
- [12] A. Morales, J. Fierrez, R. Tolosana, J. Ortega-Garcia, J. Galbally, M. Gomez-Barrero, A. Anjos, and S. Marcel, "Keystroke biometrics ongoing competition," *IEEE Access*, vol. 4, pp. 7736–7746, 2016.
- [13] *Regulation (EU) 2016/679 of the European Parliament and of the Council*, Eur. Parliament Council Eur. Union, Eur. Union, Brussels, Belgium.
- [14] P. Delgado-Santos, G. Stragapede, R. Tolosana, R. Guest, F. Deravi, and R. Vera-Rodriguez, "A survey of privacy vulnerabilities of mobile device sensors," *ACM Comput. Surv.*, vol. 54, no. 11s, pp. 1–30, Jan. 2022.
- [15] P. Melzi, C. Rathgeb, R. Tolosana, R. Vera-Rodriguez, and C. Busch, "An overview of privacy-enhancing technologies in biometric recognition," 2022, *arXiv:2206.10465*.
- [16] G. Stragapede, P. Delgado-Santos, R. Tolosana, R. Vera-Rodriguez, R. Guest, and A. Morales, "TypeFormer: Transformers for mobile keystroke biometrics," 2023, *arXiv:2212.13075*.
- [17] E. Maiorana, H. Kalita, and P. Campisi, "Mobile keystroke dynamics for biometric recognition: An overview," *IET Biometrics*, vol. 10, no. 1, pp. 1–23, Jan. 2021.
- [18] R. Giot, M. El-Abed, and C. Rosenberger, "GREYC keystroke: A benchmark for keystroke dynamics biometric systems," in *Proc. IEEE 3rd Int. Conf. Biometrics, Theory, Appl., Syst.*, Sep. 2009, pp. 1–6.
- [19] K. S. Killourhy and R. A. Maxion, "Comparing anomaly-detection algorithms for keystroke dynamics," in *Proc. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, Jun. 2009, pp. 125–134.
- [20] J. Fierrez et al., "BiosecrID: A multimodal biometric database," *Pattern Anal. Appl.*, vol. 13, no. 2, pp. 235–246, May 2010.
- [21] M. El-Abed, M. Dafer, and R. E. Khayat, "RHU keystroke: A mobile-based benchmark for keystroke dynamics systems," in *Proc. Int. Carnahan Conf. Secur. Technol. (ICCST)*, Oct. 2014, pp. 1–4.
- [22] E. Vural, J. Huang, D. Hou, and S. Schuckers, "Shared research dataset to support development of keystroke authentication," in *Proc. IEEE Int. Joint Conf. Biometrics*, Sep. 2014, pp. 1–8.
- [23] Y. Sun, H. Ceker, and S. Upadhyaya, "Shared keystroke dataset for continuous authentication," in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Dec. 2016, pp. 1–6.
- [24] C. Murphy, J. Huang, D. Hou, and S. Schuckers, "Shared dataset on natural human-computer interaction to support continuous authentication research," in *Proc. IEEE Int. Joint Conf. Biometrics (IJCB)*, Oct. 2017, pp. 525–530.
- [25] V. Dhakal, A. M. Feit, P. O. Kristensson, and A. Oulasvirta, "Observations on typing from 136 million keystrokes," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Apr. 2018, pp. 1–12.
- [26] K. Palin, A. M. Feit, S. Kim, P. O. Kristensson, and A. Oulasvirta, "How do people type on mobile devices: Observations from a study with 37,000 volunteers," in *Proc. 21st Int. Conf. Hum.-Comput. Interact. With Mobile Devices Services*, Oct. 2019, pp. 1–12.
- [27] A. Acien, A. Morales, J. Fierrez, R. Vera-Rodriguez, and O. Delgado-Mohatar, "BeCAPTCHA: Behavioral bot detection using touchscreen and mobile sensors benchmarked on HuMldb," *Eng. Appl. Artif. Intell.*, vol. 98, Feb. 2021, Art. no. 104058.
- [28] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Comput. Surv.*, vol. 54, no. 6, pp. 1–35, Jul. 2021.
- [29] I. Serna, A. Morales, J. Fierrez, and N. Obradovich, "Sensitive loss: Improving accuracy and fairness of face representations with discrimination-aware deep learning," *Artif. Intell.*, vol. 305, Apr. 2022, Art. no. 103682.
- [30] P. Terhörst, J. N. Kolf, M. Huber, F. Kirchbuchner, N. Damer, A. M. Moreno, J. Fierrez, and A. Kuijper, "A comprehensive study on face recognition biases beyond demographics," *IEEE Trans. Technol. Soc.*, vol. 3, no. 1, pp. 16–30, Mar. 2022.
- [31] A. Morales, J. Fierrez, R. Vera-Rodriguez, and R. Tolosana, "SensitiveNets: Learning agnostic representations with application to face images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 6, pp. 2158–2164, Jun. 2021.
- [32] P. Delgado-Santos, R. Tolosana, R. Guest, R. Vera-Rodriguez, F. Deravi, and A. Morales, "GaitPrivacyON: Privacy-preserving mobile gait biometrics using unsupervised learning," *Pattern Recognit. Lett.*, vol. 161, pp. 30–37, Sep. 2022.

- [33] A. Dantcheva, P. Elia, and A. Ross, "What else does your biometric data reveal? A survey on soft biometrics," *IEEE Trans. Inf. Forensics Security*, vol. 11, no. 3, pp. 441–467, Mar. 2016.
- [34] E. Gonzalez-Sosa, J. Fierrez, R. Vera-Rodriguez, and F. Alonso-Fernandez, "Facial soft biometrics for recognition in the wild: Recent works, annotation, and COTS evaluation," *IEEE Trans. Inf. Forensics Security*, vol. 13, no. 8, pp. 2001–2014, Aug. 2018.
- [35] I. Tsimperidis, A. Arampatzis, and A. Karakos, "Keystroke dynamics features for gender recognition," *Digit. Invest.*, vol. 24, pp. 4–10, Mar. 2018.
- [36] A. A. N. Buker, G. Roffo, and A. Vinciarelli, "Type like a man! Inferring gender from keystroke dynamics in live-chats," *IEEE Intell. Syst.*, vol. 34, no. 6, pp. 53–59, Nov. 2019.
- [37] A. Pentel, "Predicting user age by keystroke dynamics," in *Artificial Intelligence and Algorithms in Intelligent Systems*, R. Silhavy, Ed. Cham, Switzerland: Springer, 2019, pp. 336–343.
- [38] I. Tsimperidis, S. Rostami, K. Wilson, and V. Katos, "User attribution through keystroke dynamics-based author age estimation," in *Selected Papers from the 12th International Networking Conference*, B. Ghita and S. Shialeas, Eds. Cham, Switzerland: Springer, 2021, pp. 47–61.
- [39] A. Pentel, "Predicting age and gender by keystroke dynamics and mouse patterns," in *Proc. Adjunct Publication 25th Conf. User Model., Adaptation Personalization (UMAP)*. New York, NY, USA: Association for Computing Machinery, 2017, pp. 381–385.
- [40] I. Tsimperidis, C. Yucel, and V. Katos, "Age and gender as cyber attribution features in keystroke dynamic-based user classification processes," *Electronics*, vol. 10, no. 7, p. 835, Mar. 2021.
- [41] A. F. M. N. H. Nahin, J. M. Alam, H. Mahmud, and K. Hasan, "Identifying emotion by keystroke dynamics and text pattern analysis," *Behav. Inf. Technol.*, vol. 33, no. 9, pp. 987–996, Sep. 2014, doi: 10.1080/0144929x.2014.907343.
- [42] I. Tsimperidis, D. Grunova, S. Roy, and L. Moussiades, "Keystroke dynamics as a language profiling tool: Identifying mother tongue of unknown Internet users," *Telecom*, vol. 4, no. 3, pp. 369–377, Jul. 2023.
- [43] G. Stragapede, P. Delgado-Santos, R. Tolosana, R. Vera-Rodriguez, R. Guest, and A. Morales, "Mobile keystroke biometrics using transformers," in *Proc. IEEE 17th Int. Conf. Autom. Face Gesture Recognit. (FG)*, Jan. 2023, pp. 1–6.
- [44] J. Fierrez-Aguilar, J. Ortega-Garcia, and J. Gonzalez-Rodriguez, "Target dependent score normalization techniques and their application to signature verification," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 35, no. 3, pp. 418–425, Aug. 2005.
- [45] J. Fierrez-Aguilar, D. Garcia-Romero, J. Ortega-Garcia, and J. Gonzalez-Rodriguez, "Adapted user-dependent multimodal biometric authentication exploiting general information," *Pattern Recognit. Lett.*, vol. 26, no. 16, pp. 2628–2639, Dec. 2005.
- [46] C. Rathgeb, P. Drozdowski, D. C. Frings, N. Damer, and C. Busch, "Demographic fairness in biometric systems: What do the experts say?" *IEEE Technol. Soc. Mag.*, vol. 41, no. 4, pp. 71–82, Dec. 2022.
- [47] *Information Technology—Vocabulary, 2022. Part 37: Biometrics*, Standard ISO/IEC 2382-37(en), International Organization for Standardization, Geneva, Switzerland, 2022.
- [48] T. de Freitas Pereira and S. Marcel, "Fairness in biometrics: A figure of merit to assess biometric verification systems," *IEEE Trans. Biometrics, Behav., Identity Sci.*, vol. 4, no. 1, pp. 19–29, Jan. 2022.
- [49] P. Grother, M. Ngan, and K. Hanaoka, "Face recognition vendor test (FVRT): Part 3, demographic effects," Nat. Inst. Standards Technol., Gaithersburg, MD, USA, Tech. Rep. 8280, 2019.
- [50] J. J. Howard, E. J. Laird, Y. B. Sirotnin, R. E. Rubin, J. L. Tipton, and A. R. Vemury, "Evaluating proposed fairness models for face recognition algorithms," 2022, *arXiv:2203.05051*.
- [51] M. Abadi et al. (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. [Online]. Available: <https://www.tensorflow.org>
- [52] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," 2019, *arXiv:1912.01703*.
- [53] A. Morales, J. Fierrez, A. Acien, R. Tolosana, and I. Serna, "Setmargin loss applied to deep keystroke biometrics with circle packing interpretation," *Pattern Recognit.*, vol. 122, Feb. 2022, Art. no. 108283.
- [54] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 2, Jun. 2006, pp. 1735–1742.
- [55] S. Roy, J. Pradhan, A. Kumar, D. R. D. Adhikary, U. Roy, D. Sinha, and R. K. Pal, "A systematic literature review on latest keystroke dynamics based models," *IEEE Access*, vol. 10, pp. 92192–92236, 2022.
- [56] J. Fierrez, A. Morales, R. Vera-Rodriguez, and D. Camacho, "Multiple classifiers in biometrics. Part 2: Trends and challenges," *Inf. Fusion*, vol. 44, pp. 103–112, Nov. 2018.
- [57] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 815–823.
- [58] P. R. Borji, K. Raja, and P. Bours, "Online grooming detection: A comprehensive survey of child exploitation in chat logs," *Knowl.-Based Syst.*, vol. 259, Jan. 2023, Art. no. 110039.
- [59] M. D. Dunlop, W. Stuerzlinger, A. S. Arif, A. Komninos, and O. Alharbi, "TEXT2030-shaping text entry research in 2030," in *Proc. Adjunct Publication 24th Int. Conf. Hum.-Comput. Interact. With Mobile Devices Services*, Sep. 2022, pp. 1–4.



GIUSEPPE STRAGAPEDE received the M.Sc. degree in electronic engineering from Politecnico di Bari, Italy, in 2019. He is currently pursuing the Ph.D. degree with the Biometrics and Data Pattern Analytics (BiDA) Laboratory, Universidad Autónoma de Madrid, Spain, with a Marie Curie Fellowship, PriMa (Privacy Matters) EU Project. After one year as a Computer Vision Engineer in the industry, in 2020, he started the Ph.D. degree. His research interests include biometrics (especially mobile biometrics), data protection, signal processing, and machine learning.



RUBEN VERA-RODRIGUEZ received the Ph.D. degree in electrical and electronic engineering from Swansea University, U.K., in 2010. Since 2010, he has been with the Biometric Recognition Group, Universidad Autónoma de Madrid, Spain, where he has been an Associate Professor, since 2018. His research interests include signal and image processing, AI fundamentals and applications, HCI, forensics, and biometrics for security and human behavior analysis. He is actively involved in several National and European projects focused on these topics. He has received a Medal in the Young Researcher Awards 2022 by the Spanish Royal Academy of Engineering among other awards, and he has been a member of ELLIS Society, since 2023.



RUBEN TOLOSANA received the Ph.D. degree in computer and telecommunication engineering from Universidad Autónoma de Madrid, in 2014 and 2019, respectively. In 2014, he joined the Biometrics and Data Pattern Analytics (BiDA) Laboratory, Universidad Autónoma de Madrid, where he is currently an Assistant Professor. He is a member of the ELLIS Society, a Technical Area Committee of EURASIP, and an Editorial Board Member of the IEEE Biometrics Council

Newsletter. His research interests include signal and image processing, pattern recognition, and machine learning, particularly in the areas of deepfakes, human–computer interaction, biometrics, and health. He has also received several awards, such as the European Biometrics Industry Award, in 2018, from the European Association for Biometrics (EAB), and the Best Ph.D. Thesis Award, from 2019 to 2022, from the Spanish Association for Pattern Recognition and Image Analysis (AERFAD).



NASER DAMER (Member, IEEE) is currently a Senior Researcher with Fraunhofer IGD and a Principal Investigator and a Research Area Coordinator with the National Research Center for Applied Cybersecurity ATHENE. He is also a Lecturer with TU Darmstadt, represents the German Institute for Standardization (DIN), ISO/IEC SC37 Standardization Committee, and a member of the IEEE Biometrics Council serving on its Technical Activities Committee. He serves as an

Associate Editor for *Pattern Recognition* (Elsevier), the *Visual Computer* (Springer), and IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY.



JULIAN FIERREZ (Member, IEEE) received the Ph.D. degree in telecommunications engineering from Universidad Politécnica de Madrid, Spain, in 2006. He is currently a Full Professor with Universidad Autónoma de Madrid. His research interests include general signal and image processing, pattern recognition, and biometrics. He has been actively involved in multiple EU projects focused on biometrics, has attracted notable impact for his research, and he is a recipient of a number of distinctions. He is a member of ELLIS Society.



AYTHAMI MORALES received the Ph.D. degree from La Universidad de Las Palmas de Gran Canaria, in 2011. He has been an Associate Professor with Universidad Autónoma de Madrid, since 2017. His research interests include pattern recognition, computer vision, machine learning, and biometrics signal processing. He has participated in national and EU projects in collaboration with other universities and private entities, such as UAM, UPM, EUPMt, Indra, Union Fenosa,

Soluziona, or Accenture. He is a member of ELLIS Society.



JAVIER ORTEGA-GARCIA (Fellow, IEEE) received the Ph.D. degree (cum laude) in electrical engineering from Universidad Politécnica de Madrid, in 1996. He is currently a Full Professor with Universidad Autónoma de Madrid. He is also the Founder and the Director of the Biometrics and Data Pattern Analytics (BiDA) Laboratory. His research interests include biometric pattern recognition (on-line signature verification, speaker recognition, and human-device interaction) for security, e-health, and user profiling applications.

...