

# Deep Eyedentification: Biometric Identification using Micro-Movements of the Eye

Lena A. Jäger<sup>1</sup>(✉), Silvia Makowski<sup>1</sup>, Paul Prasse<sup>1</sup>, Sascha Liehr<sup>2</sup>, Maximilian Seidler<sup>1</sup>, and Tobias Scheffer<sup>1</sup>

<sup>1</sup> Department of Computer Science, University of Potsdam, Potsdam, Germany  
{lejaeger, silvia.makowski, prasse, maseidler, scheffer}@uni-potsdam.de

<sup>2</sup> Independent researcher  
sascha.liehr@gmail.com

**Abstract.** We study involuntary micro-movements of the eye for biometric identification. While prior studies extract lower-frequency macro-movements from the output of video-based eye-tracking systems and engineer explicit features of these macro-movements, we develop a deep convolutional architecture that processes the raw eye-tracking signal. Compared to prior work, the network attains a lower error rate by one order of magnitude and is faster by two orders of magnitude: it identifies users accurately within seconds.

**Keywords:** Machine learning · Eye-tracking · Eye movements · Deep learning · Biometrics · Ocular micro-movements

## 1 Introduction

Human eye movements are driven by a highly-complex interplay between voluntary and involuntary processes related to oculomotor control, high-level vision, cognition, and attention. Psychologists distinguish three types of macroscopic eye movements. Visual input is obtained during *fixations* of around 250 ms. *Saccades* are fast relocation movements of typically 30 to 80 ms between fixations during which visual uptake is suppressed. When tracking a moving target, the eye performs a *smooth pursuit* [21].

A large body of psychological evidence shows that these macroscopic eye movements are highly individual. For example, a large-scale study with over 1,000 participants showed that the individual characteristics of eye movements are highly reliable and, importantly, persist across experimental sessions [3]. Motivated by these findings, macro-movements of the eye have been studied for biometric identification [4,24]. Since macroscopic eye movements occur at a low frequency, long sequences must be observed before movement patterns give away the viewer’s identity—a recent study finds that users can be identified reliably after reading around 10 lines of text [33]. For use cases such as access control, this process is too slow by one to two orders of magnitude.

During fixations, the eye additionally performs involuntary micro-movements which prevent the gradual fading of the image that would otherwise occur as the

neurons become desensitized to a constant light stimulus [46,12]. *Microsaccades* have a duration ranging from 6 to 30 ms [34,35,36]. Between microsaccades, a very slow *drift* away from the center of the fixation occurs, which is superimposed by a low-amplitude, high-frequency *tremor* of approximately 40-100 Hz [34]. There is evidence that microsaccades exhibit measurable individual differences [41], but it is still unclear to what extent drift and tremor vary between individuals [28].

Video-based eye-tracking systems measure gaze angles at a rate of up to 2,000 Hz. Since the amplitudes of the smallest micro-movements are in the order of the precision of widely-used systems, the micro-movement information in the output signal is superimposed by a considerable level of noise. It is established practice in psychological research to smooth the raw eye-tracking signal, and to extract the specific types of movements under investigation. Criteria that are applied for the distinction of specific micro-movements are to some degree arbitrary [40,39], and their detection is less reliable [28]. Without exception, prior work on biometric identification only extracts macro-movements from the eye-tracking signal and defines explicit features such as distributional features of fixation durations and saccade amplitudes.

The additional information in the high-frequency and lower-amplitude micro-movements motivates us to explore the raw eye-tracking signal for a potentially much faster biometric identification. To this end, we develop a deep convolutional neural network architecture that is able to process this signal. One key challenge lies in the vastly different scales of velocities of micro- and macro-movements.

The remainder of this paper is structured as follows. Section 2 reviews prior work. Section 3 lays out the problem setting and Section 4 develops a neural-network architecture for biometric identification based on a combination of micro- and macro-movements of the eye. Section 5 presents experimental results. Section 6 discusses micro-movement-based identification in the context of other biometric technologies; Section 7 concludes.

## 2 Related Work

There is a substantial body of research on biometric identification using macro-movements of the eye. Most work uses the same stimulus for training and testing—such as a static cross [4], a jumping point [22,23,44,8,47], a face image [43,17,5], or various other kinds of images [10]. Using the same known stimulus for training and testing opens the methods to replay attacks.

Kinnunen and colleagues present the first approach that uses different stimuli for training and testing and does not involve a secondary task; they identify subjects who watch a movie sequence [27]. Later approaches use eye movements on novel text to identify readers [19,30].

A number of methods have been benchmarked in challenges [25,24]. All participants in these challenges and all follow-up work [45] present methods that extract saccades and fixations, and define a variety of features on these macro-movements, including distributions of fixation durations and of amplitudes, ve-

locities, and directions of saccades. Landwehr and colleagues define a generative graphical model of saccades and fixations [30] from which Makowski and colleagues derive a Fisher Kernel [33]; Abdelwahab *et al.* develop a semi-parametric discriminative model [2]. All known methods are designed to operate on an eye-gaze sequence of considerable length; for example, one minute of watching a video or reading about one page of text.

### 3 Problem Setting

We study three variations of the problem of biometric identification based on a sequence  $\langle (x_0, y_0), \dots, (x_n, y_n) \rangle$  of yaw gaze angles  $x_i$  and pitch gaze angles  $y_i$  measured by an eye tracker. For comparison with prior work, we adopt a *multi-class classification* setting. For each user from a fixed population of users, one or more enrollment eye-gaze sequences are available that are recorded while the user is reading text documents. A multi-class classification model trained on these enrollment sequences recognizes users from this population at application time while the users are reading different text documents. Classification accuracy serves as performance metric in this setting.

Multi-class classification is a slight abstraction of the realistic use case in two regards. First, this setting disregards the possibility of encountering a user from outside the training population of users. Secondly, the learning algorithm has to train the model on enrollment sequences of all users. This training would have to be carried out on an end device or a cloud backend whenever a new user is enrolled; this is unfavorable from a product perspective.

In the more realistic settings of *identification* and *verification*, an embedding is trained offline on eye-gaze sequences for training stimuli of a population of training identities. At application time, the model encounters users from a different population who may view different stimuli. Users are enrolled by simply storing the embedding of their enrollment sequences. The model identifies a user when a similarity metric between an observed sequence and one of the enrollment sequences exceeds a recognition threshold.

In the *identification* setting, multiple users can be enrolled. Since the ratio of enrolled users to impostors encountered by the system at application time is not known, the system performance has to be characterized by two ROC curves. One curve characterizes the behavior for enrolled users; here, false positives are enrolled users who are mistaken for different enrolled users. The second curve characterizes the behavior for impostors; false positives are impostors who are mistaken for one of the enrolled users.

In the *verification* setting, the model verifies a user’s presumed identity. This setting is a special case of identification in which a single user is enrolled. As no confusion of enrolled users is possible, a single ROC curve characterizes the system performance.

## 4 Network Architecture

We transform each eye-gaze sequence  $\langle(x_0, y_0), \dots, (x_n, y_n)\rangle$  of absolute angles into a sequence  $\langle(\delta_1^x, \delta_1^y), \dots, (\delta_n^x, \delta_n^y)\rangle$  of angular gaze velocities in  $^\circ/\text{s}$  with  $\delta_i^x = r(x_i - x_{i-1})$  and  $\delta_i^y = r(y_i - y_{i-1})$ , where  $r$  is the sampling rate of the eye tracker in Hz.

The angular velocity of eye movements differs greatly between the different types of movement. While drift occurs at an average speed of around  $0.1\text{-}0.4^\circ/\text{s}$  and tremor at up to  $0.3^\circ/\text{s}$ , microsaccades move at a rapid  $15$  to  $120^\circ/\text{s}$  and saccades even at up to  $500^\circ/\text{s}$  [34,36,40,21]; there is, however, no general agreement about the exact cut-off values between movement types. Global normalization of the velocities squashes the velocities of drift and tremor to near-zero and models trained on such data resort to extracting patterns only from macro-movements. For this reason, our key design element of the architecture consists of independent subnets for *slow* and *fast* movements which observe the same input sequences but with different scaling.

Both subnets have the same number and type of layers; Figure 1 shows the architecture. Both subnets process the same sliding window of 1,000 velocity pairs which corresponds to one second of input data, but the input is scaled differently. Equation 1 transforms the input such that the low velocities that occur during tremor and drift roughly populate the value range between  $-0.5$  and  $+0.5$  while velocities of microsaccades and saccades are squashed to values between  $-0.5$  and  $-1$  or  $+0.5$  and  $+1$ , depending on their direction. The parameter  $c$  has been tuned within the range of psychologically plausible values from  $0.01$  to  $0.06$ .

$$t_s(\delta_i^x, \delta_i^y) = (\tanh(c\delta_i^x), \tanh(c\delta_i^y)) \quad (1)$$

Equation 2, in which  $z(\cdot)$  is the  $z$ -score normalization, truncates absolute velocities that are below the minimal velocity  $\nu_{min}$  of microsaccades and saccades. Based on the psychological literature, the threshold  $\nu_{min}$  was tuned within the range of  $10$  to  $60^\circ/\text{s}$ .

$$t_f(\delta_i^x, \delta_i^y) = \begin{cases} z(0) & \text{if } \sqrt{\delta_i^{x^2} + \delta_i^{y^2}} < \nu_{min} \\ (z(\delta_i^x), z(\delta_i^y)) & \text{otherwise} \end{cases} \quad (2)$$

Each subnet consists of 9 pairs of one-dimensional convolutional and average-pooling layers. The model performs a batch normalization on the output of each convolutional layer before applying a ReLU activation function and performing average pooling. Subsequently, the data feeds into two fully connected layers with batch-normalization and ReLU activation with a fixed number of  $2^8$  and  $2^7$  units, followed by a fully connected layer of  $2^7$  units with ReLU activation that serves as embedding layer for identification and verification. For classification and for the purpose of training the network in the identification and verification setting, this is followed by a softmax output layer with a number of units equal to the number of training identities that is discarded after training in the identification and verification settings.

**Table 1.** Parameter space used for grid search: kernel size  $k$  and number of filters  $f$  of the convolutional layers, the scaling parameter  $c$  of Equation 1 and the velocity threshold  $\nu_{min}$  of Equation 2.

Parameter	Search space
$c$	{0.01, 0.02, 0.04, 0.06}
$\nu_{min}$	{10°/s, 20°/s, 30°/s, 40°/s, 60°/s}
$k$	{3, 5, 7, 9}
$f$	{32, 64, 128, 256, 512}

**Table 2.** Best hyperparameter configuration found via grid search in the search space shown in Table 1.

Parameter	Layer	Slow subnet	Fast subnet
$c$	$t_s$	0.02	-
$\nu_{min}$	$t_f$	-	40°/s
$k$	conv 1-3	9	9
	conv 4-7	5	5
	conv 8-9	3	3
$f$	conv 1-3	128	32
	conv 4-7	256	512
	conv 8-9	256	512

Figure 1 shows the overall architecture which we refer to as the *DeepEyedentification* network. The output of the subnets is concatenated and flows through a fully connected layer of  $2^8$  units and a fully connected layer with  $2^7$  units that serves as embedding layer for identification and verification, both with batch normalization and ReLU activation. The overall architecture is trained in three steps. The fast and the slow subnets are pre-trained independently and their weights are frozen in the final step where the joint architecture is trained.

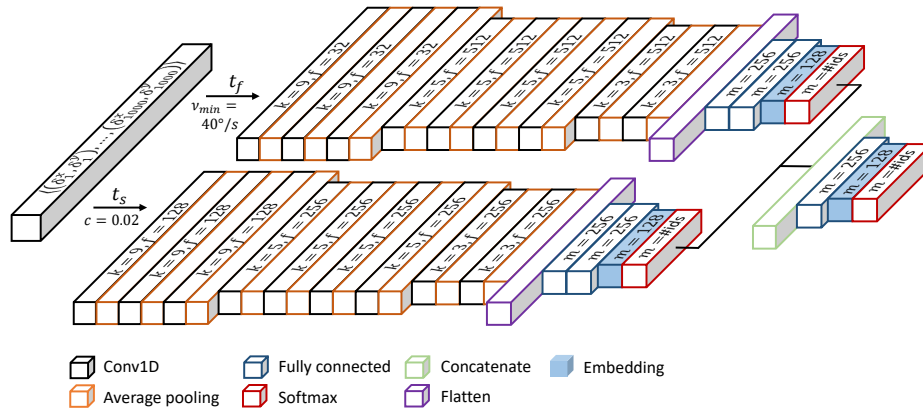
In the identification and verification settings, the final embedding consists of the concatenation of the joint embedding and the embeddings generated by the fast and slow subnets. In this case, the cosine similarity serves as metric for the comparison of enrollment and input sequences.

## 5 Experiments

This section reports on experiments in the settings of multi-class classification, identification, and verification. All code is available at <https://osf.io/ps9qj/>.

### 5.1 Data Collection

We use two data collections for our experiments. Makowski *et al.* [33] have collected the largest eye-tracking data set for which the raw output signal is available. It consists of monocular eye-tracking data sampled at 1,000 Hz from 75 participants who are reading 12 scientific texts of approximately 160 words. In



**Fig. 1.** Network architecture. Parameter  $c$  denotes the scaling factor of Equation 1,  $v_{min}$  the velocity threshold of Equation 2,  $k$  the kernel size,  $f$  the number of filters and  $m$  the number of fully connected units. Batch normalization and ReLU activation are applied to the output of all convolutional and fully connected layers. All convolutional layers have a stride of 1; all pooling layers have a pooling size of 2 and a stride of 1.

order to extract absolute gaze angles, the eye tracker has to be calibrated for each participant. Makowski *et al.* exclude data from 13 participants whose data is poorly calibrated. Since DeepEyedentification only processes velocities, we do not exclude any data. We refer to this data set as *Potsdam Textbook Corpus*.

The Potsdam Textbook Corpus was acquired in a single session per user. To explore whether individuals can be recognized across sessions, we collect an additional data set from 10 participants à four sessions with a temporal lag of at least one week between any two sessions. We record participants’ gaze using a binocular Eyelink Portable Duo eye tracker at a sampling rate of 1,000 Hz. During each session, participants are presented with 144 trials in which a black point consecutively appears at 5 random positions on a light gray background on a  $38 \times 30$  cm monitor ( $1280 \times 1024$  px). The interval in which the point changes its location varies between trials (250, 500, 1000 or 1500 ms). We refer to these data as *JuDo* (Jumping Dots) data set. We use the Potsdam Textbook Corpus for hyperparameter optimization, and evaluation of the DeepEyedentification network in a multi-class classification and an identification and verification setting, while we use the much smaller JuDo data set to assess the model’s session bias.

## 5.2 Reference Methods

Existing methods for biometric identification using eye movements only operate on macroscopic eye movements; they first preprocess the data into sequences of saccades and fixations and use different features computed from these macro-movements such as fixation duration or saccade amplitude. Existing methods

that allow different stimuli for training and testing can be classified into i) approaches which aggregate the extracted features over the relevant recording window, ii) statistical approaches that compute the similarity of scanpaths by applying statistical tests to the distributions of the extracted features, and iii) graphical models that generate sequences of fixation durations and saccade amplitudes. As representative aggregational reference method, we choose the model by Holland and Komogortsev (2011) that is specifically designed for eye movements in reading [19]. As statistical reference approaches we use the first model of this kind by Holland and Komogortsev (2013) [20] and the current state-of-the-art model by Rigas *et al.* (2016) [45]. As representative graphical models, we also use the first model of this kind by Landwehr *et al.* (2014) [30] and the state-of-the-art model by Makowski *et al.* (2018) [33].

### 5.3 Hyperparameter Tuning

We optimize the hyperparameters via grid search on one hold-out validation text from the Potsdam Textbook Corpus which we subsequently exclude from the training and testing of the final network; Table 1 gives an overview of the space of values and Table 2 the selected values that we keep fixed for all subsequent experiments. We vary the kernel sizes and numbers of filters of each subnet independently, but constrain them to be identical within convolutional layers 1-3, 4-7, and 8-9. Moreover, we constrain the kernel size to be smaller or equal and the number of filters to be greater or equal compared to the preceding block.

### 5.4 Hardware and Framework


We train the networks on a server with a 40-core Intel(R) Xeon(R) CPU E5-2640 processor and 128 GB of memory and a GeForce GTX TITAN X GPU using the NVidia CUDA platform with Tensorflow version 1.12.0 [1] and Keras version 2.2.4 [7]. As optimizer, we use Adam [26,42] with a learning rate of 0.001 for the training of the subnets and 0.0001 for the common layers. All models and submodels are trained with a batch size of 64 sequences.

### 5.5 Multi-Class Classification

This section focuses on the multi-class classification setting in which the model is trained on the Potsdam Textbook Corpus to identify users from a fixed population of 75 users who are represented in the training data, based on an eye-gaze sequence for an unseen text.

In this setting, data are split across texts, to ensure that the same stimulus does not occur in training and test data. We perform leave-one-text-out cross validation over 11 text documents. We study the accuracy as a function of the duration of the eye-gaze signal. For each duration, we let the model process a sliding window of 1,000 time steps and average the scores over all window positions.

The reference models are evaluated on the same splits. They receive preprocessed data as input: sequences are split into saccades and fixations, and the relevant fixation and saccade features are computed. At test time, these models receive only as many macro-movements as input as fit into the duration.



**Fig. 2.** Multi-class classification on the Potsdam Textbook Corpus. Categorical accuracy as a function of the amount of available test data in seconds; error bars show the standard error. The models marked with \* are evaluated on a subset of the data containing 62 well-calibrated users, all other methods are evaluated on the full data set of 75 readers.

Figure 2 shows that for any duration of an input sequence, the error rate of DeepEyedentification is roughly one order of magnitude below the error rate of the reference methods. DeepEyedentification exceeds an accuracy of 91.4% after one second, 99.77% after 10 seconds and reaches 99.86% accuracy after 40 seconds of input, whereas Rigas *et al.* [45] reach 8.37% accuracy after one second and 43.02% after 10 seconds, and the method of Makowski *et al.* [33] reaches 91.53% accuracy after 100 seconds of input. We can conclude that micro-movements convey substantially more information than lower-frequency macro-movements of the eye.


The figure also shows that the overall network is significantly more accurate than either of its subnets. The *fast subnet*, for which only velocities of microsaccades and saccades are visible while tremor and drift are truncated to zero, reaches an accuracy of approximately 77% after one second. The *slow subnet*, which perceives the velocities of tremor and drift on an almost-linear scale while the velocities of microsaccades and saccades are squashed by the sigmoidal transformation, achieves roughly 88% of accuracy after one second.

## 5.6 Identification and Verification

In these settings, the input window slides over the test sequence and an enrolled user is identified (true positive) if and when the cosine similarity between any



input window and any window in his enrollment sequence exceeds the recognition threshold; otherwise, the user counts as a false negative. A false positive arises when the similarity between a test sequence from an enrolled user (confusion setting) or an impostor (impostor setting) and the enrollment sequence of a different user exceeds the threshold; otherwise a true negative arises. We perform 50 iterations of random resampling on the Potsdam Textbook Corpus. In each iteration, we randomly draw 50 training users and train the DeepEyedentification model on 9 training documents for these users. One text serves as enrollment sequence and one text remains as observation. In the identification setting, a randomly drawn set of 20 of the 25 users who are not used for training are enrolled, and the remaining 5 users act as impostors. In the verification setting, one user is enrolled and 24 impostors remain.




(a) Confusions between 20 enrolled users. (b) Confusions between an unknown number of impostors and 20 enrolled users.


**Fig. 3.** Identification on the Potsdam Textbook Corpus. ROC curves for the confusion setting (a) and the impostor setting (b) as a function of the duration of the input signal at application time, both with 20 enrolled users. Error bars show the standard error.

For the identification setting, Figure 3a shows the ROC curves for confusions between the 20 enrolled users on a logarithmic scale. The area under the ROC curve increases from 0.9787 for one second of data to 0.9915 for 10 and 0.9964 after 90 seconds; the corresponding EER values are 0.09, 0.04, and 0.02. Figure 3b shows the ROC curves for confusions between an impostor and one of the 20 enrolled users; here, the AUC values lie between 0.7382 and 0.9385, the corresponding EER values between 0.31 and 0.1.


Figure 4 shows the ROC curve for the verification setting. Here, the AUC lies between 0.9556 for one second, 0.9879 for 10, and 0.9924 for 90 seconds. In this setting, each impostor can only be confused with one presumed identity, whereas, in the identification setting, an impostor can be confused with each of the 20 enrolled users. Figure 5 shows a *t*-SNE visualization [31] that illustrates how



**Fig. 4.** Verification on the Potsdam Textbook Corpus. ROC curves for the confusions between one enrolled user and an unknown number of impostors as a function of the duration of the input signal at application time. Error bars show the standard error.



**Fig. 5.** *t*-sne visualization of the embedding for 10 users.

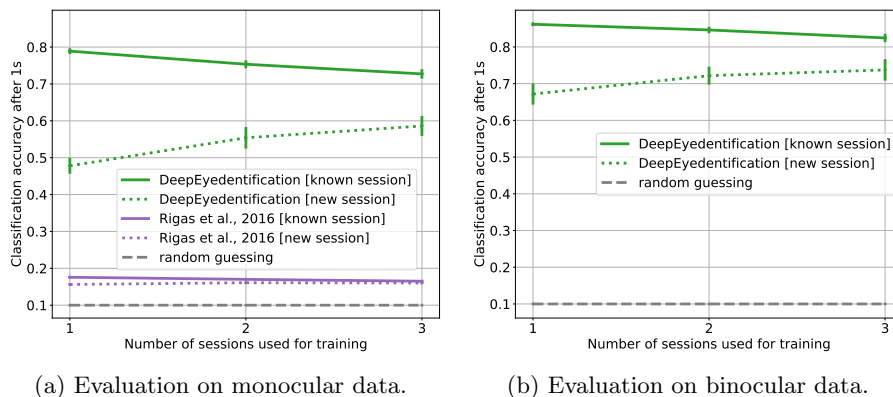


**Fig. 6.** Time to classification with standard error over false-positive rate.

the embedding layer clusters 10 users randomly drawn from outside the training identities. Finally, Figure 6 shows the time to identification as a function of the false-positive rate for the identification and verification settings.

### 5.7 Assessing Session Bias

Using the JuDo data set, we investigate the DeepEyedidentification network’s ability to generalize across recording sessions by comparing its multi-class classification performance on test data taken either from the same sessions that are used for training or from a new session. We train the DeepEyedidentification network and the reference method that performed best on the Potsdam Textbook Corpus [45] on one to three sessions using the same hyperparameters and learning framework as for the main experiments (see Sections 5.3 and 5.4). We



**Fig. 7.** Multi-class classification on the JuDo data set. Categorical accuracy on one second of test data from either a known or a new recording session as a function of the number of sessions used for training with a constant total amount of training data. The results are averaged over ten iterations for each held-out test session. Error bars show the standard error.

evaluate the models using leave-one-session-out cross-validation on one held-out session (test on a new session) and on 20% held-out test data from the remaining session(s) (test on a known session). When training on multiple sessions, the amount of training data from each session is reduced such that the total amount of data used for training remains constant. Since binocular data is available, we also evaluate the DeepEyedentification network on binocular data by applying it independently to synchronous data from both eyes and averaging the softmax scores of the output layer. At training, the data from the two eyes are treated as separate instances.

Figure 7a shows the results for monocular test sequences of one second. After one second of input data, the model reaches a classification accuracy of 78.9% when testing and training it on data from a single session, and an accuracy of up to 58.6% when training and testing it on different sessions. Increasing the number of training sessions reduces the session bias significantly ( $p < 0.05$  for one versus three sessions). The model of Rigas *et al.* [45] reaches accuracies around 16% in all settings. The use of binocular data (see Figure 7b) not only improves the overall performance of the DeepEyedentification network, but also significantly reduces the session bias compared to monocular data ( $p < 0.05$  for one training session). When being trained on three sessions, the model achieves an accuracy of 73.75% on a new test session after only one second of input data.

## 5.8 Additional Exploratory Experiments

We briefly summarize the outcome of additional exploratory experiments. First, we explore the behavior of a variant of the DeepEyedentification architecture that has only a single subnet which processes the globally normalized input.

This model does not exceed the performance of the fast subnet, which indicates that it extracts only macro-movement patterns.

Second, we find that adding an input channel that indicates whether a time step is part of a fixation or part of a saccade according to established psychological criteria [14,15] does not improve the model performance. Moreover, forcing the slow subnet to only process movements during fixations and forcing the fast subnet to only process movements during saccades deteriorates the model performance. Our interpretation of this finding is that given the amount of information contained in the training data, an established heuristic categorization of movement types contributes no additional value.

Lastly, we change the convolutional architecture into a recursive architecture with varying numbers of LSTM units [18]. We find that the convolutional architecture consistently outperforms the explored LSTM architectures.

## 6 Discussion

This section discusses eye movements in relation to other biometric technologies. We discuss relevant qualitative properties of biometric methods: the required level of user interaction, the population for which the method can be applied, attack vectors, and anti-spoofing techniques.

While fingerprints and hand-vein scans require an explicit user action—placing the finger or the hand on a scanning device—face identification, scanning the iris, and tracking micro-movements of the eye can in principle be performed unobtrusively, without explicit user interaction. Scanning the iris or recording the micro-movements of the eye without requesting the user to step close to a camera would, however, require a camera that offers a sufficiently high resolution over a sufficiently wide field of view.

Biometric technologies differ with respect to intrinsic limitations of their applicability. For instance, fingerprints are worn down by hard physical labor, iris scanning requires users with small eyes to open their eyes unnaturally wide and is not available for users who wear cosmetic contact lenses. Since micro-movements of the eye are a prerequisite for vision, this method applies to a large potential user base.

All biometric identification methods can be attacked by acquiring biometric features from an authorized user and replaying the recorded data to the sensor. For instance, face identification can be attacked by photographs, video recordings, and 3D masks [16]. A replay attack on ocular micro-movement-based identification is theoretically possible but requires a playback device that is able to display a video sequence in the infrared spectrum at a rate of 1,000 frames per second. Biometry can similarly be attacked by replaying recorded or artificially generated data during enrollment. For instance, wearing cosmetic contact lenses during enrollment with an iris scanner can cause the scanner to accept other individuals who wear the same contact lens as false positives [37].

Anti-spoofing techniques for all biometric technologies firstly aim at detecting imperfections in replayed data; for example, missing variation in the input

over time can indicate a photograph attack. This problem is intrinsically difficult because it is an adversarial problem; an attacker can always minimize artifacts in the replayed data. As an illustration, an attacker can replay a video recording instead of a still image to add liveliness. Liveliness detection is implicitly included in identification based on eye movements. Secondly, additional sensors can be added—such as multi-spectral cameras or depth sensors to prevent photograph-based and video-based replay attacks. This of course comes at additional costs and can still be attacked with additional effort, such as by using 3D-printed models instead of photographs. Thirdly, the identification procedure can include a randomized challenge to which the user has to respond. For example, a user can be asked to look at specific positions on a screen [32,29,11,13,48,9]. Challenges prevent replay attacks at the cost of obtrusiveness, bypassing them requires a data generator that is able to generate the biometric feature and also respond to the challenge. Identification based on movements of the eye is unique: responding to challenges demands neither the user’s attention nor a conscious response. Randomized salient stimuli in the field of view immediately trigger an involuntary eye movement that can be validated.

## 7 Conclusion

Our research adds to the list of machine-learning problems for which processing raw input data with a deep CNN greatly improves the performance over methods that extract engineered features. In this case, the improvement is particularly remarkable and moves a novel biometric-identification technology close to practical applicability. The error rate of the DeepEyedentification network is lower by one order of magnitude and identification is faster by two orders of magnitude compared to the best-performing previously-known method.

We would like to point out that at this point the embedding layer of DeepEyedentification has been trained with 50 users. Nevertheless, it attains a true-positive rate of 60% at a false-positive rate of 1% after two seconds of input in the verification setting. By comparison, the embedding layer of a current face-identification model that attains a true-positive rate of 95.6% at a false-positive rate of 1% has been trained with 9,000 users [6]. A recent iris-recognition model attains a true-positive rate of 83.8% at a false-positive rate of 1% [38]. This comparison highlights the high potential of identification based on micro-movements.

We have developed an approach to processing input that contains signals on vastly different amplitudes. Global normalization squashes the velocities of the most informative, high-frequency but low-amplitude micro-movements to nearly zero, and networks which we train on this type of input do not exceed the performance of the fast subnet. The DeepEyedentification network contains two separately trained subnets that process the same signal scaled such that the velocities of slow movements, in case of the slow subnet, and of fast movements, in case of the fast subnet, populate the input value range.

Biometric identification based on eye movements has many possible fields of application. In contrast to fingerprints and hand-vein scans, it is unobtrusive.

While iris scans fail for cosmetic contact lenses and frequently fail for users with small eyes, it can be applied for all individuals with vision. A replay attack would require a device able to display 1,000 frames per second in the infrared spectrum. Moreover, replay attacks can be prevented by including a challenge in the form of a visual stimulus in the identification procedure to which the user responds with an involuntary eye movement without assigning attention to the task.

## Acknowledgments

This work was partially funded by the German Science Foundation under grant SFB1294, and by the German Federal Ministry of Research and Education under grant 16DIII116-DII. We thank Shravan Vasishth for his support with the data collection.

## References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems. <https://www.tensorflow.org/> (2015)
2. Abdelwahab, A., Kliegl, R., Landwehr, N.: A semiparametric model for Bayesian reader identification. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP 2016). pp. 585–594 (2016)
3. Bargary, G., Bosten, J.M., Goodbourn, P.T., Lawrance-Owen, A.J., Hogg, R.E., Mollon, J.: Individual differences in human eye movements: An oculomotor signature? *Vision Research* **141**, 157–169 (2017)
4. Bednarik, R., Kinnunen, T., Mihaila, A., Fränti, P.: Eye-movements as a biometric. In: Proceedings of the 14th Scandinavian Conference on Image Analysis (SCIA 2005). pp. 780–789 (2005)
5. Cantoni, V., Galdi, C., Nappi, M., Porta, M., Riccio, D.: GANT: Gaze analysis technique for human identification. *Pattern Recognition* **48**, 1027–1038 (2015)
6. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: VGGFace2: A dataset for recognising faces across pose and age. In: 13th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2018). pp. 67–74 (2018)
7. Chollet, F., et al.: Keras. <https://keras.io> (2015)
8. Cuong, N., Dinh, V., Ho, L.S.T.: Mel-frequency cepstral coefficients for eye movement identification. In: 24th International Conference on Tools with Artificial Intelligence (ICTAI 2012). pp. 253–260 (2012)
9. Cymek, D., Venjakob, A., Ruff, S., Lutz, O.M., Hofmann, S., Roetting, M.: Entering PIN codes by smooth pursuit eye movements. *Journal of Eye Movement Research* **7**, 1–11 (2014)
10. Darwish, A., Pasquier, M.: Biometric identification using the dynamic features of the eyes. In: 6th International Conference on Biometrics: Theory, Applications and Systems (BTAS 2013). pp. 1–6 (2013)

11. De Luca, A., Weiss, R., Hußmann, H., An, X.: Eyepass – eye-stroke authentication for public terminals. In: *Extended Abstracts on Human Factors in Computing Systems (CHI EA '08)*. pp. 3003–3008 (2007)
12. Ditchburn, R.W., Ginsborg, B.L.: Vision with a stabilized retinal image. *Nature* **170**, 36–37 (1952)
13. Dunphy, P., Fitch, A., Olivier, P.: Gaze-contingent passwords at the ATM. In: *4th Conference on Communication by Gaze Interaction (COGAIN)*. pp. 59–62 (2008)
14. Engbert, R., Kliegl, R.: Microsaccades uncover the orientation of covert attention. *Vision Research* **43**, 1035–1045 (2003)
15. Engbert, R., Mergenthaler, K.: Microsaccades are triggered by low retinal image slip. *Proceedings of the National Academy of Sciences of the USA* **103**, 7192–7197 (2006)
16. Erdogmus, N., Marcel, S.: Spoofing face recognition with 3D masks. *IEEE Transactions on Information Forensics and Security* **9**(7), 1084–1097 (2014)
17. Galdi, C., Nappi, M., Riccio, D., Cantoni, V., Porta, M.: A new gaze analysis based softbiometric. In: *5th Mexican Conference on Pattern Recognition (MCPR 2013)*. pp. 136–144 (2013)
18. Graves, A., Schmidhuber, J.: Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks* **18**(5-6), 602–610 (2005)
19. Holland, C., Komogortsev, O.V.: Biometric identification via eye movement scanpaths in reading. In: *2011 International Joint Conference on Biometrics (IJCB 2011)*. pp. 1–8 (2011)
20. Holland, C., Komogortsev, O.: Complex eye movement pattern biometrics: Analyzing fixations and saccades. In: *2013 International Conference on Biometrics (ICB-2013)* (2013)
21. Holmqvist, K., Nyström, M., Andersson, R., Dewhurst, R., Jarodzka, H., Van de Weijer, J.: *Eye tracking: A comprehensive guide to methods and measures*. Oxford University Press, Oxford (2011)
22. Kasprowski, P.: *Human identification using eye movements*. Ph.D. thesis, Silesian University of Technology, Poland (2004)
23. Kasprowski, P., Ober, J.: Enhancing eye-movement-based biometric identification method by using voting classifiers. In: *Proceedings of SPIE 5779: Biometric Technology for Human Identification II*. pp. 314–323 (2005)
24. Kasprowski, P., Harkeżlak, K.: The second eye movements verification and identification competition. In: *Proceedings of the International Joint Conference on Biometrics* (2014)
25. Kasprowski, P., Komogortsev, O.V., Karpov, A.: First eye movement verification and identification competition at BTAS 2012. In: *Proceedings of the IEEE Fifth International Conference on Biometrics: Theory, Applications and Systems (BTAS 2012)*. pp. 195–202 (2012)
26. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
27. Kinnunen, T., Sedlak, F., Bednarik, R.: Towards task-independent person authentication using eye movement signals. In: *Proceedings of the 2010 Symposium on Eye-Tracking Research and Applications (ETRA '10)*. pp. 187–190 (2010)
28. Ko, H.K., Snodderly, D.M., Poletti, M.: Eye movements between saccades: Measuring ocular drift and tremor. *Vision Research* **122**, 93–104 (2016)
29. Kumar, M., Garfinkel, T., Boneh, D., Winograd, T.: Reducing shoulder-surfing by using gaze-based password entry. In: *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS 2007)*. pp. 13–19 (2007)

30. Landwehr, N., Arzt, S., Scheffer, T., Kliegl, R.: A model of individual differences in gaze control during reading. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*. pp. 1810–1815 (2014)
31. Maaten, L.v.d., Hinton, G.: Visualizing data using t-SNE. *Journal of Machine Learning Research* **9**, 2579–2605 (2008)
32. Maeder, A., Fookes, C., Sridharan, S.: Gaze based user authentication for personal computer applications. In: *Proceedings of the 2004 International Symposium on Intelligent Multimedia, Video and Speech Processing*. pp. 727–730 (2004)
33. Makowski, S., Jäger, L.A., Abdelwahab, A., Landwehr, N., Scheffer, T.: A discriminative model for identifying readers and assessing text comprehension from eye movements. In: *Machine Learning and Knowledge Discovery in Databases (ECML PKDD 2018)*. pp. 209–225 (2019)
34. Martinez-Conde, S., Macknik, S.L., Hubel, D.H.: The role of fixational eye movements in visual perception. *Nature Reviews Neuroscience* **5**, 229–240 (2004)
35. Martinez-Conde, S., Macknik, S.L., Troncoso, X.G., Dyar, T.A.: Microsaccades counteract visual fading during fixation. *Neuron* **49**, 297–305 (2006)
36. Martinez-Conde, S., Macknik, S.L., Troncoso, X.G., Hubel, D.H.: Microsaccades: A neurophysiological analysis. *Trends in Neurosciences* **32**, 463–475 (2009)
37. Morales, A., Fierrez, J., Galbally, J., Gomez-Barrero, M.: Introduction to iris presentation attack detection. In: *Handbook of Biometric Anti-Spoofing*, pp. 135–150. Springer (2019)
38. Nalla, P.R., Kumar, A.: Toward more accurate iris recognition using cross-spectral matching. *IEEE Transactions on Image Processing* **26**, 208–221 (2017)
39. Nyström, M., Hansen, D.W., Andersson, R., Hooge, I.: Why have microsaccades become larger? Investigating eye deformations and detection algorithms. *Vision Research* **118**, 17–24 (2016)
40. Otero-Millan, J., Troncoso, X.G., Macknik, S.L., Serrano-Pedraza, I., Martinez-Conde, S.: Saccades and microsaccades during visual fixation, exploration, and search: Foundations for a common saccadic generator. *Journal of Vision* **8**(14), 21–21 (2008)
41. Poynter, W., Barber, M., Inman, J., Wiggins, C.: Individuals exhibit idiosyncratic eye-movement behavior profiles across tasks. *Vision Research* **89**, 32–38 (2013)
42. Reddi, S.J., Kale, S., Kumar, S.: On the convergence of Adam and beyond. In: *International Conference on Learning Representations (ICLR 2018)* (2018)
43. Rigas, I., Economou, G., Fotopoulos, S.: Biometric identification based on the eye movements and graph matching techniques. *Pattern Recognition Letters* **33**, 786–792 (2012)
44. Rigas, I., Economou, G., Fotopoulos, S.: Human eye movements as a trait for biometrical identification. In: *Fifth International Conference on Biometrics: Theory, Applications and Systems (BTAS 2012)*. pp. 217–222 (2012)
45. Rigas, I., Komogortsev, O., Shadmehr, R.: Biometric recognition via eye movements: Saccadic vigor and acceleration cues. *ACM Transactions on Applied Perception* **13**(2), 6 (2016)
46. Riggs, L.A., Ratliff, F.: The effects of counteracting the normal movements of the eye. *Journal of the Optical Society of America* **42**, 872–873 (1952)
47. Srivastava, N., Agrawal, U., Roy, S., Tiwary, U.S.: Human identification using linear multiclass SVM and eye movement biometrics. In: *8th International Conference on Contemporary Computing (IC3)*. pp. 365–369 (2015)
48. Weaver, J., Mock, K., Hoanca, B.: Gaze-based password authentication through automatic clustering of gaze points. In: *2011 IEEE International Conference on Systems, Man, and Cybernetics (SMC 2011)*. pp. 2749–2754 (2011)