Using Reflexive Eye Movements for Fast Challenge-Response Authentication

Ivo Sluganovic, Marc Roeschlin, Kasper B. Rasmussen, Ivan Martinovic

Department of Computer Science, University of Oxford
first.last@cs.ox.ac.uk

ABSTRACT

Eye tracking devices have recently become increasingly popular as an interface between people and consumer-grade electronic devices. Due to the fact that human eyes are fast, responsive, and carry information unique to an individual, analyzing person's gaze is particularly attractive for effortless biometric authentication. Unfortunately, previous proposals for gaze-based authentication systems either suffer from high error rates, or require long authentication times.

We build upon the fact that some eye movements can be reflexively and predictably triggered, and develop an interactive visual stimulus for elicitation of reflexive eye movements that supports the extraction of reliable biometric features in a matter of seconds, without requiring any memorization or cognitive effort on the part of the user. As an important benefit, our stimulus can be made unique for every authentication attempt and thus incorporated in a challenge-response biometric authentication system. This allows us to prevent replay attacks, which are possibly the most applicable attack vectors against biometric authentication.

Using a gaze tracking device, we build a prototype of our system and perform a series of systematic user experiments with 30 participants from the general public. We investigate the performance and security guarantees under several different attack scenarios and show that our system surpasses existing gaze-based authentication methods both in achieved equal error rates (6.3%) and significantly lower authentication times (5 seconds).

1. INTRODUCTION

Eye tracking devices capture precise position and movement of the human cornea on a millisecond scale. This in turn allows determining the exact location of one's gaze on a screen or on surrounding objects. Since analyzing eye behavior can give insight into our internal cognitive processes and even predict conditions such as autism [24], eye trackers have been used in neurophysiological research for over a century, but until recently their use in everyday life was limited due to prohibitive equipment costs.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CCS'16 October 24-28, 2016, Vienna, Austria

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4139-4/16/10.

DOI: <http://dx.doi.org/10.1145/2976749.2978311>

However, the speed and responsiveness of eye movements strongly motivate their use as an attractive input channel for human-computer interaction; as a result, recent years have brought a sharp reduction in retail prices of eye tracking devices. While dedicated trackers can be purchased for as little as \$100 [1], eye tracking capabilities are also being added to consumer products such as laptops [23], cars [34], tablets, and mobile phones [32]. Given the diverse advantages and applications of eye tracking, its widespread expansion into our everyday lives is only likely to continue.

As we demonstrate in the following sections, tracking a user's gaze is particularly suitable for fast and low-effort user authentication, especially in scenarios where keyboard input is not available. Eye movements exhibit traits distinctive enough that classification algorithms (e.g., [13]) can reliably discern among a large group of individuals. However, despite the advantages, exploiting eye movements for user authentication remains a challenging topic. As we summarize in Section 8, previous work on gaze-based authentication achieves either high error rates (e.g., EER above 15%) or long authentication times (e.g., above 20 seconds). One likely explanation for some of these outcomes are overly complex visual stimuli that result in voluntarily triggered eye movements which are highly dependent on a user's current cognitive state.

In this paper, we show how the reflexive physiological behavior of human eyes can be used to build fast and reliable biometric authentication systems. We utilize the fact that, even though most eye movements are elicited voluntarily, specific *reflexive* movements can be actively triggered using a simple visual stimulus. Measuring and analyzing millisecond-scale characteristics of reflexive eye movements provides several important benefits. Users' eyes naturally and spontaneously react to the shown stimulus so they do not need to follow any instructions or memorize additional information. As a result, elicitation of reflexive behavior requires lower cognitive load and is very fast. This in turn enables keeping authentication times short while at the same time extracting large amounts of useful biometric data and achieving low error rates.

Finally, we show a crucial advantage of exploiting reflexive eye movements for authentication: by employing a challenge-response type of protocol, such systems can provide security even under a stronger adversary model than is usually considered for biometrics. One of the obstacles for widespread use of biometric authentication in our daily lives is the fact that most biometrics can be captured and replayed relatively easily. Examples include spoofing image

recognition systems with photographs from social media and spoofing fingerprint recognition using copies of fingerprints left behind on everyday items. If the visual stimulus can be made unique for each authentication attempt, then the elicited responses will accordingly be different, but still include user-specific characteristics. By always choosing a new *challenge* (randomly generated stimulus) and verifying if the *response* (measured eye movements) corresponds to it, our authentication system can assert that the biometric sample is indeed fresh. Other biometric systems have to make special provisions to achieve a level of spoofing and replay protection. For example, sophisticated fingerprint readers measure additional attributes like temperature and moisture in order to determine liveness. Our gaze-based authentication system obtains these guarantees practically for free, without requiring any other information besides the recording of a user's eye movements.

2. BACKGROUND ON EYE MOVEMENTS

We start with a short background of the human visual system and necessary eye movements terminology; this allows us to introduce main concepts that motivate our research and guide the design of the system in the following sections.

Even when one's gaze is firmly fixated on a single stimulus, human eyes are never completely still. They are constantly making hundreds of micro movements per second, which are interlaced with more than 100,000 larger movements during the course of one day [2]. During visual tasks, such as search or scene perception, our eyes alternate between *fixations* and *saccades*. Fixations are used to maintain the visual focus on a single stimulus, while saccades reorient the eye to focus the gaze on a next desired position. Saccades are rapid eye movements and they are considered to be the fastest rotational movement of any external part of our body, reaching angular velocities of up to 900 degrees per second, and usually lasting between 20 ms and 100 ms. In Figure 1, fixations can be seen as areas of large numbers of closely grouped points, while saccades consist of series of more spread recordings that depict fairly straight paths.

When a salient change happens in our field of vision, our eyes naturally reorient on the target, since this is a necessary first step to provide information for further higher-level cognitive processes [30]. These externally elicited saccades happen reflexively and are considered to be an effortless neuronal response, requiring very low cognitive load from the user. After the stimulus onset, a corresponding *reflexive* saccade is initiated rapidly, with usual latencies of less than 250 ms. In contrast, voluntary saccadic movements have larger mean latencies (above 300 ms) which are additionally influenced by different internal and external factors [41].

The analysis of eye movements has been part of medical research for more than a century since it offers valuable information of our cognitive and visual processing [30, 9, 3]. Keeping the goal of reliable biometric authentication in mind, we are interested in extracting and combining multiple characteristics of human eye movements for which there exists supporting research that they offer stable individual differences between users. For example, Castelhamo et al. [8] examine stable individual differences in characteristics of both saccades and fixations and provides support for their stable use in biometric authentication. Saccades were also used in [13] to enable stable authentication and identification. Furthermore, several researchers have analyzed

Figure 1: Eye movements of four different users as a response to the same visual stimulus. Fixations are visible as clustered areas, while saccades consist of series of dots that depict paths. Larger red dots represent the positions at which the visual stimulus was shown. Despite their distinct characteristics, all four gaze paths closely match the positions of the stimulus.

eye behavior features of trained shooters [12], professional baseball players [4] and other specific groups of individuals [15], and reported measurable differences between their eye movements characteristics.

Given that reflexive reactions are less dependent on momentary conscious states of an individual than conscious actions, it is expected that biometrics based on reflexive characteristics offer more stable authentication. Furthermore, taking into account the advantage in faster elicitation times, the goal of our research is to design a stimulus that supports the use of reflexive saccades for biometric authentication. For example, prior research has shown that saccade latencies depend on the dominant eye [31, 26] of the individual, which is a stable characteristic and provides strong motivation for using saccade latencies for classification. Finally, it was shown that *saccade latency* varies if anticipation (temporal expectancy) is present [40]. This provides an argument for randomizing the stimulus that is shown to users.

3. ASSUMPTIONS AND GOALS

We start by defining the system and adversary model used throughout this paper; we then state the design goals for the visual stimulus and the authentication system.

System Model. We assume the general settings of a user authenticating to a workstation in an office scenario throughout the course of a normal work day. A simple visualization of the system model is shown in Figure 2. The user authenticates to a workstation using a gaze tracking device throughout the course of a workday. The workstation uses data acquired by the gaze tracker and a user's biometric template to make the authentication decision.

A legitimate user is one who is enrolled with the authen-

Figure 2: System model. The workstation uses data acquired by the gaze tracker and user’s biometric template to make the authentication decision. The adversary has read-write access to the gaze channel. The visual channel is authenticated and therefore read-only.

tication system. The enrollment happens in a secure scenario, where the legitimate user authenticates to the workstation using some other authentication method. During enrollment, the user is shown several visual stimuli and the workstation uses the corresponding recordings of the user’s gaze to create a biometric template used for identity verification.

The interaction takes place through three different channels. The *visual channel* is an authenticated channel from the workstation to the user that consists of a screen that displays information, and the *gaze tracking channel* from the user to the gaze tracker allows the workstation to determine characteristics about the user’s eyes, including where he is looking on the screen, as well as capture the reflexive eye movements described in Section 2.

The workstation itself cannot be modified or forced to run unintended code.

Adversary Model. The adversary’s goal is to impersonate a legitimate user and successfully authenticate to the workstation. The adversary can freely choose his victim from the set of enrolled users. Since he can observe both the visual and gaze channels, the adversary has access to the biometric data from previous authentication attempts by the victim.

We focus on two different types of attacks that the adversary can perform:

- *Impersonation attack.* The adversary tries to gain access to the workstation by positioning himself in front of the gaze tracking device. This is the most common way of evaluating biometric authentication systems, and is usually reported in terms of false reject (FRR) and false accept rates (FAR) as well as equal error rates (EER).
- *Replay attack.* The adversary targets a specific user and replays his previously recorded authentication attempt to the authentication system. This can be done either at the sensor level (e.g. by using a mechanical eye replica), or by bypassing the gaze tracking sensor completely and injecting the recorded samples between the workstation and the sensor.

Biometrics are non-revocable, and we are surrounded by sensors that can be used to steal and replay biometric data. Therefore, we believe that modeling an attacker as having access to legitimate user’s previous biometric measurements is a realistic and necessary assumption. Most static biometrics, such as fingerprints or face recognition [5], cannot provide security under such assumptions; the ability to prevent

replay attacks is one of the major strengths of our scheme since simply replaying an acquired sample is arguably the most accessible attack vector for most biometrics.

We do not consider a targeted adversary who is able to model and generate arbitrary artificial samples of a user’s eye movements in an interactive manner. As we further discuss in Section 9, such attacks require significantly higher levels of complexity and effort from the adversary; a level of commitment against which most biometric systems can not provide security guarantees.

Design Goals.

- *Low cognitive load:* The system should pose low cognitive load on the users. Ideally, users should not be required to remember credentials, carry tokens, or learn new procedures. Moreover, the cooperation required from the user should be as effortless as possible.
- *Fast:* The duration of a single authentication attempt should be as short as possible.
- *Resistance against replay:* The system should make it difficult for an adversary to replay acquired biometric samples and thereby successfully authenticate.

4. SYSTEM ARCHITECTURE

We propose an authentication system which works as follows. The workstation shows an interactive visual stimulus on the screen (we refer to it as *gaze-challenge*). Simultaneously, the gaze tracking device captures eye movements of the user as he watches the screen (*gaze-response*), which the workstation uses to adapt the stimulus in real time. Finally, the workstation makes a decision about the user’s identity and verifies if the received gaze-response corresponds to the shown gaze-challenge, asserting that the captured eye movements are indeed fresh.

4.1 Stimulus for Reflexive Saccade Elicitation

To achieve stated design goals, a visual stimulus should satisfy several requirements. It should elicit responses that are sufficiently distinctive to allow discrimination between different users. The response should not require high cognitive effort and should not depend on a user’s momentary cognitive state. The stimulus should be *unpredictable* to prevent habituation: seeing an image for the first time will likely result in a different response than seeing it for the second and the consecutive times [40]. Finally, in order to allow fast authentication, the stimulus duration should be as short as possible.

Design. Considering that reflexive behavior is more stable and less dependent on a user’s transient internal cognitive states than voluntary behavior, our goal is to design a stimulus which allows eliciting and measuring individual traits of user’s reflexive saccadic responses. Reflexive saccades are triggered by salient objects that appear in one’s field of view; thus our stimulus consists of presenting a single red dot on a dark screen that changes position multiple times. As shown in Figure 3, a user’s eyes respond to the change by eliciting a reflexive saccade which reorients the gaze towards the dot. Every time the position of the dot changes, the visual system responds by initiating a new reflexive saccade. Due to saccade latency, this happens after a period of 100-200 ms during which the visual system processes new information.

Figure 3: A visualization of the stimulus for reflexive saccade elicitation. At any given time, only a single red dot is shown; previous positions are shown on this figure to help the reader. Shortly after a red dot appears on the screen (a), a user’s visual system starts a reflexive saccade to shift the gaze (dotted path) towards its position. Several milliseconds later, as the user’s gaze enters the invisible perimeter around the stimulus (dashed circles), the dot is considered successfully gazed and momentarily changes its position. Before a new saccade starts, there is usually a fixation lasting 100-250 ms, during which the visual system processes new input information (saccade latency). In (d), the presented dot is again successfully gazed, and once more changes its position.

Figure 4: Relative frequency of saccade latencies for gaze-responses used in this paper. Latencies are computed as the duration between the stimulus change and the start of subsequent saccadic movement. Vertical lines discriminate between reflexive and other types of saccades based on [41]; latencies of reflexive saccades are usually lower than 250 ms, in contrast to latencies of voluntary saccades that are over 250 ms. Values under 80 ms are likely the result of noise or blinks, or voluntary saccades initiated well before the stimulus change.

Ideally, our stimulus should elicit the maximal number of reflexive saccades in a given period of time, and this highly depends on the frequency with which the position of the dot changes. If this frequency is too high, user’s eyes will not be given sufficient time to perform a full saccade. If it is too low, the user might get tired of looking at a static screen and start voluntary saccadic movements. Furthermore, each user is slightly different, so there might not exist a unique frequency at all. Our stimulus ensures an optimum between these trade-offs by interactively changing the location of the dot as soon as the user *successfully gazes* the dot, i.e., when a user’s gaze enters a perimeter of radius r around the dot’s center. This results in eliciting the maximal number of full saccades in any given time interval, and ensures that the user’s visual system receives an outside stimulus change as often as possible, thus reducing the elicitation of voluntary saccades which depend on his current cognitive state. To ensure that the stimulus terminates even if the user is not looking at the screen, the dot is considered to be unsuccessfully gazed and moves to the next position after a specific period of $t_{\max}$ milliseconds has passed. This process continues for all N stimulus positions that constitute a gaze-challenge.

Basing an authentication system on reflexive movements

provides additional benefits: taking into account that reflexive behavior is significantly harder to consciously control, an adversary is less likely to be able to successfully imitate another user’s characteristics. Most importantly, because of the natural and effortless tendency of the human visual system to keep “catching” the red dot, the response to such visual stimulus is fully reflexive: users neither need to follow specific instructions nor invest high cognitive effort —*their eyes do the work themselves*.

Effectiveness of the stimulus. In order to evaluate how effectively our designed stimulus elicits reflexive behavior, we compute saccade latencies for a total of 991 gaze measurements that constitute the dataset used throughout this paper. Since each of the measurements represents a gaze-response to a stimulus with 25 different positions for the dot, in total, this sums up to analyzing close to 25,000 captured saccades.

Figure 4 shows the distribution and categorization of the measured saccade latencies, dividing them into reflexive saccades, voluntary saccades and saccadic movement caused by blinks. Latencies under 80 ms are physically impossible and likely to be the result of blinks or noise. Remaining latencies predominantly fall below 250 ms, the threshold that characterizes reflexive saccades [41]. This lets us conclude that the stimulus does indeed elicit primarily reflexive behavior.

4.2 Authentication Protocol

We now use the proposed stimulus as a building block in a challenge-response protocol for biometric user authentication that is secure against replay attacks. At the end of the protocol execution, the workstation knows if the user whose identity is claimed is at the moment present in front of the gaze tracking device. To that goal, the workstation must ensure that two properties hold:

Freshness. Freshness of the received biometric data can be ensured by always showing a different randomly generated visual stimulus (gaze-challenge) to which every response will differ in a verifiable way.

Correct Identity. The user has the ability to generate biometric data that corresponds to the claimed user’s template which was created during enrollment.

The protocol for local biometric authentication is shown in Figure 5. After the user claims his identity, the workstation

Figure 5: Biometric challenge-response authentication protocol. User claims his identity, after which the workstation generates a fresh *gaze-challenge* c that is an ordered list of positions in which the stimulus is shown. User looks at (**LookAt**) a screen where the stimulus is shown at N positions $\{(x_i, y_i)\}$. Meanwhile, the gaze tracking device records the user's gaze paths g_i for all stimulus positions that constitute the *gaze-response* $\vec{g}$. The workstation verifies the freshness of $\vec{g}$, and finally verifies that the biometric features extracted from $\vec{g}$ correspond to the claimed identity.

generates a fresh visual stimulus, which we refer to as *gaze-challenge* (c_W) in the rest of the paper. c_W consists of a set of n randomly chosen coordinates, which uniquely define the interactive stimulus described in Section 4.1. As the gaze-challenge is presented to the user, his eyes reflexively respond with a series of eye movements, which constitute the *gaze-response* (r_U). Gaze-response is recorded by the gaze tracking device through the gaze channel.

In order to accept or reject the user's authentication request, the workstation performs two verification steps: **VerifyFreshness** and **VerifyIdentity**. These are described in detail in Sections 4.3 and 4.4, respectively.

In the final message, the workstation notifies the user if he has been granted or denied access to the system.

4.3 VerifyFreshness

As described in Section 4.1, each visual stimulus is uniquely defined by a list of N coordinates; therefore, it is possible to always present a different random gaze-challenge to the user. Since no visual stimulus shown to users is ever reused, in order to verify the freshness of the response, it suffices to verify if the received gaze-response closely corresponds to the freshly presented gaze-challenge. As visualized in Figure 3, if some gaze-response was recorded while specific gaze-challenge was shown to the user, then the user's eye movements should closely resemble the order and positions in which the stimulus dot was shown. This is visible in Figure 1: despite differences in gaze patterns of different users, all of them correspond to the locations of the stimulus dot.

The system determines if the gaze-response is indeed fresh by ensuring that the user timely gazed at the majority of the stimulus positions. After a stimulus dot is shown in

one of the N positions, it is considered *successfully gazed* only if one of the subsequent measurements of the user's gaze position falls within a radius of r pixels from the center of the stimulus dot. Otherwise, if no gaze measurement falls within its radius after $t_{\max}$ milliseconds, a position is considered to be unsuccessfully gazed and the dot moves to the next position:

$$g_i := [(x_j, y_j) : t_i \leq t_j < t_i + t_{\max}]$$

$$\text{gazed}(x_i, y_i) \iff \exists (x, y) \in g_i : \|(x, y) - (x_i, y_i)\|_2 \leq r$$

In order to decide on the freshness of the received gaze-response, the system checks if the ratio of successfully gazed stimulus positions is greater or equal to a chosen percentage threshold T .

As the threshold T increases, the possibility that an adversary successfully replays an old recording of a legitimate user's gaze decreases. On the other hand, this also results in more legitimate attempts failing freshness verification, e.g., because of imprecise gaze measurements. We explore this trade-off and the security guarantees provided by our system against replay attacks in Section 7.3.

4.4 VerifyIdentity

If the received gaze-response passes the freshness verification, the system finally verifies that it truly originated from the user whose identity was claimed at the beginning of the authentication. The received gaze-response is first used as input to compute a set of specific feature values that are idiosyncratic and support stable classification between users. Next, the computed features are used as an input to a two-class classifier which is created during user enrolment. The classifier determines whether the calculated features more likely belong to the user whose identity was claimed, or to some internal or external attacker. As a last step, the authentication system makes a final decision and notifies the user of acceptance or rejection.

Next section describes the details about the features that we use and how we train the user classifiers.

5. GAZE CLASSIFICATION

This section describes the process of extracting individual characteristics from a user's gaze-response and training a classifier that can uniquely discriminate between future responses of the same user and any other user's gaze patterns.

5.1 Saccade and Fixation Detection

The first step to computing feature values is to split the input gaze measurement into intervals of saccades and fixations.

We implement an algorithm [33] that estimates the level of noise in the data and adaptively determines the thresholds which are used to classify the pattern into periods of fixations and saccades based on angular speeds and accelerations. As seen in Figure 6a, the algorithm also detects eye movement recordings that could not have been generated by a human eye under known physiological constraints, and are usually the result of blinking. Given that the mean duration of a single blink is close to 200 ms [18], and that head movements and gazes outside of the screen area usually last even longer, it is important to denoise the raw data before further analysis. These artefacts are filtered based on research that shows the peak angular speed of the human

Figure 6: Visualization of features on (a) temporal and (b) spatial plots of a the raw gaze tracking data. In Subfigure (a), the moment when stimulus changes position is depicted with a vertical red line. The period depicted with horizontal stripes is physiologically impossible for a human eye to perform and is caused by a blink. We remove such artefacts with methods described in Section 5.

eye to lie between 700 and 900 *deg/sec* [18], and the peak angular acceleration to not cross 100000 *deg/sec²*.

Having grouped the measurements as belonging either to a fixation or a saccade, we proceed to calculate a set of features for each recorded gaze sample, ignoring those measurements that are classified as noise by the procedure.

5.2 Feature Extraction

We next compute the features for gaze classification. The features should be as varied for different users and as similar as possible for multiple authentication attempts of the same user.

As Figure 6a shows, each gaze-response consists of intermixed periods of saccades and fixations, and each such period allows us to compute multiple features. However, we are interested in computing a set of identifiable feature values for the whole gaze-response, irrespective of the number of elicited saccades and fixations; to that end, and to reduce the effect of noise, feature values for a single gaze-response are computed as the median of feature values computed on individual saccades or fixations in that gaze measurement.

Since not all potential features contribute the same amount of distinguishing power, we follow a semi-automated approach to select the optimal set of features for the authentication system. Initially, we explore a broader set of fixation and saccade traits, in addition to a range of other metrics that measure overall characteristics of the gaze path. Based on the Relative Mutual Information (RMI), we test the features on randomly chosen subsets of the data set, measure their classification performance, and exclude those that do not achieve satisfactory results. The RMI values of the resulting features that we use in the remainder of the paper can be found in Table 1, while Figure 6 illustrates the extraction procedure.

As the RMI values in Table 1 show, medians of **average angular speeds** during fixations or saccades, as well as the **duration** of fixations are among the most specific features we tested. This finding is congruent with the feature assessment conducted by Eberz et al. [13], where pairwise speeds exhibit the highest relative mutual information, only outperformed by some of their static features, such as pupil

diameter. Contrary to their results, we identify **saccade curviness** (ratio of air distance and total distance of a saccade) and **saccade latency** to be the features that yield the most distinguishing power. Furthermore, we identify several discriminative features based on computing a convex hull of all measurements in a fixation: **convex hull and circumference**, as well as **fixation density**, defined as the ratio of the convex hull area and the number of gaze measurements in that fixation.

This paper only uses dynamic characteristics of eye movements; we thus purposely forego using several potentially discriminating features that the gaze tracking devices can provide, such as an estimate of user’s pupil size and the distances between the user’s eyes. In prior work, pupil size was shown to be a discriminative feature for gaze-based authentication systems [13], however, the authors raise valid concerns that an adversary could manipulate his pupil size, e.g., by controlling the lightning conditions. Despite potential classification improvements, in order to provide a more conservative estimate of the performance that gaze-based authentication systems can achieve, in this paper we choose to employ only features that can be extracted from raw co-ordinates of the user’s gaze.

5.3 User Enrolment

During enrolment, several gaze-responses are used to train a dedicated 2-class classifier that the system will use as user’s identity verifier: based on the set of feature values extracted from any subsequent gaze-response, the classifier makes a decision whether the values correspond to him or not.

Besides legitimate user’s gaze-responses, the enrolment procedure requires a similarly sized set of gaze-responses belonging to other users that are labeled as negative samples during classifier training.

We use a Support Vector Machine (SVM) [10] with Radial Basis Function (RBF) kernel as the classifier, since SVMs are known to provide strong classification results for non-linear data sets. We also evaluated other classification algorithms on a subset of the data, and confirmed that SVMs achieved stronger classification than other evaluated statistical models (Random Forrest and AdaBoost Trees).

Table 1: Relative Mutual Information (RMI) of an assortment of the most informative features

Median of saccade	RMI	Median of fixation	RMI	Overall	RMI
Duration	0.1864	Duration	0.1959	Avg. time per stimulus	0.1824
Avg. speed	0.1921	Avg. speed	0.2150	Avg. distance per stimulus	0.1927
Max. speed	0.1709	Max. speed	0.1968	Avg. speed	0.2053
Latency	0.2041	Max. distance to center	0.1604		
Max acceleration	0.1675	Convex hull area	0.1894		
Ratio air/total distance	0.2397	Convex hull circumference	0.1899		
		Density	0.2063		

SVMs with RBF kernels are fully defined by two hyperparameters: **1)** C , which controls the trade-off between the penalty of incorrect classification and the margin of the decision hyperplane, and **2)** σ , which is a parameter that defines the scale of the radial basis function. The optimal pair of hyper-parameter values is chosen from a predetermined set of potential values, based on the evaluation that uses 5-fold cross-validation: for each pair of potential hyperparameters, 80% of the enrolment data is used to train the resulting classifier, while the remaining 20% of the enrolment data is used to evaluate the classification performance; this is repeated five times.

The pair of hyperparameters that resulted in strongest classification performance is finally used to derive the final user classifier which is used in future authentication.

6. DATA ACQUISITION

In order to experimentally evaluate the performance of the proposed system and protocol, we developed a prototype and ran a series of user experiments to gather data for analysis.

6.1 System Prototype

Setup. Our prototype setup is composed of a gaze tracking device (SMI RED 500 [38]), a 24-inch LED screen and a desktop computer. The generation of the visual stimulus and the gaze sampling was performed by a custom-built software library that controls the gaze tracking device. We implemented procedures that take care of the internal calibration of the gaze tracker, the measurement of the sampling accuracy and the visual presentation of the stimulus, as well as the acquisition of the gaze samples captured by the gaze tracker.

Parameters. For each authentication attempt, the system generated a visual challenge consisting of $N = 25$ random dot positions. Red stimulus dot was shown on a plain dark background, with a diameter of 0.7 cm. In order to detect that a dot was successfully gazed, we used a perimeter radius of $r = 1.4$ cm. If not successfully gazed, the dot changed position after $t_{\max} = 1000$ ms. The distance between users' eyes and the gaze tracking device (positioned directly underneath the screen) was 70 cm.

6.2 User Experiments

Experiment Design. For the purpose of assessing feasibility and performance of the proposed system, we conducted a series of user experiments that reflect the scenario described in Section 3. We refer to a series of consecutive authentication attempts with the same participant as one session. Each session lasted about 10 minutes and included a briefing and 15 authentication attempts. Before participant's first session

we generated a calibration profile that was reused during all subsequent sessions with that participant. To analyze the performance of our system, both from the perspective of a user and an attacker, we divided the participants into two groups: legitimate users who have completed the enrollment procedure, and external attackers, whose gaze characteristics were not known to the system.

In order to show that our system can successfully authenticate users over the course of a normal work day (without re-calibration), we require each enrolled user to take part in a minimum of three (up to four) sessions. The first two sessions are five minutes apart and mimic a legitimate user leaving his desk to take a break or use the restroom. All subsequent sessions are at least 6 hours apart. Participants acting as external attackers are only invited to one session where they are asked to impersonate a legitimate user, i.e., the system uses the calibration profile and biometric template of the chosen legitimate user. Every external attacker tries to authenticate as 5 different legitimate users, at least 3 times per user. In their last session, legitimate users were asked to act as internal attackers and each performed a minimum of 15 attempts of impersonating other users, analogously to external attackers.

Test Population. Experimental data was acquired from a total of 30 participants aged 21 to 58 who were recruited from the general public through public advertisements, email lists, and social media. The only requirement was a minimum age of 18. The test population consists of 7 women and 23 men. Out of the 30 recruited participants, 22 participants were enrolled as legitimate users and 8 participants represented external attackers whose gaze characteristics were not known to the system. The acquired data set consists of a total of 1602 gaze-responses: 1021 authentication attempts by legitimate users and 581 simulated attack attempts by either internal or external attackers.

Participants were told that their eye movements will be recorded for the purpose of evaluating the feasibility of distinguishing individuals based on their behavioral gaze-based biometrics. They then signed a written consent form in accordance with the experiment ethics review approved by the University's research ethics committee, reference number SSD/CUREC1A/14-226. Names have been replaced with pseudonyms.

Participants who do not have normal vision wore contact lenses or were asked to remove their glasses. This was done to remove the possibility that classification relies on potential specific characteristics of recorded gaze when glasses are worn. For the same reason, lighting conditions were not changed during all experiment sessions.

Figure 7: Measured authentication time and EER as a function of gaze-challenge complexity N . As N increases from 8 to 24, the EER reduces from above 12% to under 6%, while at the same time, the median time to authenticate grows linearly from 2 seconds to about 9 seconds. The vertical line depicts a scenario where 15 positions are used in a challenge: the median authentication time is around 5 seconds, while the EER is close to 7%.

7. SYSTEM EVALUATION

We now experimentally evaluate the proposed system with respect to the design goals stated in Section 3.

7.1 Varying the Challenge Complexity N

One of the defining parameters of the proposed system is N , the number of stimulus positions in a single gaze-challenge. We first analyze the effect that varying N has on authentication time and overall user classification performance. Incrementing N directly increases the complexity of gaze-challenge, thus requiring more time to respond to the visual stimulus. At the same time, larger N should allow the system to extract more stable features and thus achieve stronger classification results. On the other hand, as N decreases, both the authentication time, and the classification performance are likely to decline.

Setup. Since all user experiments were run with gaze-challenges that had $N = 25$ stimulus dot positions, we can evaluate the classifier performance in a scenario where gaze-challenges consist of $K < N$ positions by simulating that the stimulus presentation and gaze recording stopped after the K -th position was gazed. Such an adapted dataset is constructed by only considering gaze measurements that were recorded before the $(K + 1)$ -th stimulus position is shown.

The classification performance for each K and for each user is estimated by computing an Equal Error Rate (EER) while performing a five-fold cross-validation of the individual classifiers as follows. In each of five repetitions, four out of five folds of the legitimate user’s authentication attempts are provided as enrolment data for user enrolment that was performed as described in Section 5. The remaining fold was used to evaluate classifier performance against other users’ authentication attempts as negative samples. The resulting EER for any K is computed as an average across all five folds of all individual users’ classifiers for that K .

Results. We show the effect of varying N on authentication time and classification performance in Figure 7. The median time for a single authentication attempt grows linearly from

Figure 8: Empirical cumulative distribution function for duration of all measured authentication attempts when $N = 15$. Close to 50% of the attempts took less than 5 seconds, while more than 80% of the attempts lasted less than 7.5 seconds.

2 seconds for 8 stimulus positions, to about 9 seconds for 24 stimulus positions. At the same time, the overall EER of the classification falls from around 12% when only 8 stimulus positions are used, to a level of 6% when 24 stimulus positions are used in a challenge.

Since $N = 15$ shows a balanced trade-off between classification performance and median authentication time, we use this value to report results in the remainder of the analysis. In order to provide a more comprehensive estimate of the time required for the majority of users to authenticate than just median, in Figure 8 we show a cumulative density function of the authentication times for all users when $N = 15$. The figure shows that half of the users authenticate in 5 seconds or less, while the authentication for more than 80% of the users takes less than 7.5 seconds. As we discuss in Section 9, these times are favorable to previous related work in gaze-based authentication, as well as reported password authentication times.

7.2 Impersonation Attacks

Setup. Recall that, in an impersonation attack, the attacker targets a specific user with the goal of responding to the gaze-challenge posed by the system, and successfully impersonating the legitimate user in order to gain access. The attacker is permitted to use the gaze-based authentication system in any way he wishes, such as purposely moving or altering the angle of his head to try to increase the chance of gaining access.

As described in Section 6.2, we purposely design the user experiments to simulate this type of attack as closely as possible: all participants were asked to perform multiple “attack attempts”, in which they falsely claimed some other user’s identity and tried to authenticate with the gaze calibration profile of the legitimate user loaded by the system.

For each user, we perform a five-fold cross-validation to estimate the performance of the system under such attacks. We enrol the user as described in Section 5, using four out of five folds of legitimate user’s samples, and then evaluate the performance of the whole authentication system on the remaining one fifth of the legitimate user’s gaze-responses that were not used for enrolment. During evaluation, legitimate user’s samples are labeled as positive, while all attack attempts that other users made while pretending to be the legitimate user are labeled as negative. We consider an authentication attempt accepted by the system only if it passes both the identity verification and the freshness verification. For freshness verification, we use a threshold $T = 50\%$.

Figure 9: The ROC curves that show authentication performance under impersonation attacks. Red and green curves represent only internal and external attackers, while blue curve shows the overall combined performance. The EER for internal attackers equals to 6.2%, while for external attackers it is expectedly slightly higher, and amounts to 7.3%. The overall EER for all attackers is 6.3%.

Besides overall performance, we also separately evaluate two disjunct subsets of the attack attempts: those originating from external attackers, who are unknown to the system, and those originating from internal attackers, whose previous authentication attempts might have been used as negative samples during enrollment.

Results. We show the system performance against impersonation attacks as an ROC curve in Figure 9. Since individual user classifiers output a probability that a given sample belongs to the respective legitimate user, we can achieve different classification performance by varying the threshold above which a sample is considered legitimate. As this threshold increases, so does the likelihood of falsely rejecting a legitimate user (FRR) increase, but at the same time, the likelihood of falsely accepting an attacker (FAR) decreases. Different combinations of FAR and FRR values for three attack scenarios (internal, external, and all attackers) are shown in Figure 9. For all three scenarios, it is possible to achieve low FAR values (under 5%) if FRR is increased closer to 10% and vice-versa.

An Equal Error Rate (EER) is defined as the rate at which FRR and FAR are equal, and is usually used to compare different classifiers. As expected, in terms of EER, the system achieves slightly stronger performance against internal attackers (6.2% EER) than external attackers (7.3% EER). Overall, the system achieves an EER of 6.3% for impersonation attacks; as we discuss in Section 8, this result is preferable to any previously reported performance of gaze-based authentication systems.

7.3 Replay Attacks

Setup. Recall from Section 4.3 that in order to prevent reuse of biometric data, the system verifies that the received gaze-response corresponds to the presented gaze-challenge, i.e., that the user successfully gazed at no less than a chosen percentage T of the stimulus positions presented during authentication.

The result of verifying freshness of a received response

Figure 10: Performance of the freshness verification procedure depending on the chosen threshold T . As we change the required percentage of successfully gazed stimuli to classify a gaze sample as “fresh” from 0% to 50%, the ratio of successfully detected replay attempts rises from 0 to close to 1. At the same time, the ratio of successfully classified fresh attempts starts declining as the required threshold increases over 60%, showing almost perfect results for the thresholds between 40% and 60%.

does not depend on the claimed identity during authentication, but only on the positions of the dot in the visual stimulus. Therefore, in order to provide a more comprehensive estimate of the distinctiveness of a challenge-response pair, we report the results for a scenario in which identity verification always returns a positive answer.

In order to evaluate the probability of success of a replay attack, for each gaze-challenge c^i , we simulate a “replay” of all other gaze-responses g^j to the `VerifyFreshness` function of the system. We calculate the success rate of replaying g^j to c^i as the percentage of stimulus positions from c^i that would be considered successfully gazed if a user’s response was g^j .

Since our dataset consists of 1021 legitimate authentication attempts, each recorded with a unique gaze-challenge, we are able to simulate more than 10^6 potential replay attempts in order to estimate the true reject rate. Furthermore, in order to estimate the true accept rates, we use the same procedure to simulate a total of 1021 legitimate authentication attempts, in which the gaze-response was indeed generated as the user was presented the matching gaze-challenge.

Results. Figure 10 shows achieved performance of the challenge-response verification for different values of T , which we vary from 0% to 100%. As T , the ratio of replay attempts that are correctly rejected (TRR) increases, while the ratio of legitimate, fresh attempts that are correctly accepted (TAR) decreases.

A desired threshold is the one that detects all replay attempts, while accepting all legitimate authentication attempts as fresh. Figure 10 shows a wide range of potential threshold values that lie between 40% and 60% and almost perfectly separate the fresh and the replayed gaze-responses. Such a broad range of thresholds that achieve strong classification is a desirable property for any classification system as it gives strong confidence in reported results.

Since we use $T = 50\%$ to evaluate impersonation attacks, we report specific numeric details for this threshold. The results of simulating more than 10^6 challenge-response pairs as replay attempts show that we achieve close to perfect true reject rates (TRR) of 99.94%. At the same time, very few legitimate attempts are incorrectly rejected: the evaluation shows a true accept rate (TAR) of 98.63%, a result of falsely rejecting only 14 out of 1021 legitimate attempts.

Table 2: Comparison to existing biometric authentication systems based on eye-movements

	Analysis of	Stimulus	Ref.	Time [s]	EER [%]	Notes
high-level f.	Scan paths + arch densities	Human faces	[7]	17	25	
	Distribution of areas of interest	Human faces	[14]	10	36.1	
	Graph matching	Human faces	[35]	4	30	
	Fixation density maps	Movie trailer	[36]	60	14	
low-level f.	Cepstrum transform of raw signal	Dot, fixed inter-stimulus	[22]	8	N/A	FAR 2%, FRR 22%
	Oculomotor plant model	Dot, horizontal sequence	[27]	21	N/A	FAR 5.4%, FRR 56.6%
	Scan paths and fixation features	Read section of text	[17]	60	23	
	Fixation and saccade features	Read section of text	[16]	60	16.5	
	Liveness detection	Dot, horizontal sequence	[28]	100	18	Focus on liveness detection
	Fixation features and pupil sizes	Click the dot	[13]	40	7.8	Continuous authentication
	Fixation and saccade features	Dot, interactive	this paper	5	6.3	Replay: FAR 0.06%

Overall, these results show that our system robustly prevents replay attempts for a wide range of thresholds with very high success rates. Furthermore, taking into account that the system can detect repeated replay attempts, and e.g. lock user’s account after certain number of failed attempts, we finally conclude that our system can effectively prevent replay attacks.

8. RELATED WORK

While different eye tracking methods have been used in medical research for over a century, their use in security is fairly recent. A review paper by Zhang et. al. [43] provides an overview of authentication methods and systems proposed before 2010, while Saeed [37] gives a more recent comparison of methods and results of gaze-based authentication systems proposed up to year 2013. According to Zhang et. al. [43], existing work in user identification and authentication can be roughly divided into two categories: **1)** using gaze tracking as a human-computer interface (control channel) to support standard security primitives and **2)** using characteristics of the gaze patterns to extract individual biometric traits that enable distinguishing between different users.

In the first line of research, individuals use their eyes to prove their identity by naturally and covertly inputting secret information such as passwords [29, 42, 6] or specific patterns on the screen [5, 11, 25]. Using eyes as a control channel has several advantages, such as prevention of shoulder-surfing and smudge attacks. Unfortunately, these approaches usually share the negative characteristics of passwords, such as requiring the users to learn a procedure or remember and recall different pieces of information, as well as still being susceptible to eavesdropping and replay attacks.

Our work belongs to the second, biometric approach that uses the characteristics of individual’s gaze patterns to discriminate between different users. Such authentication systems usually come with the general benefits, but also challenges typical of biometrics: they usually require no memorization, prevent sharing of credentials, and offer high usability, but at the same time, they suffer from irrevocability, which renders replay attacks a serious threat if even a single user’s biometric sample is acquired by an attacker.

Biometric approaches to gaze-based authentication can be further divided into two subcategories: those that rely on high-level characteristics of user’s gaze patterns (*where* and *what* the user is looking at), and those that analyze the low-level traits of *how* the user’s eyes are moving.

High-level Characteristics. The first approach is motivated by hypotheses that users exhibit individual behavior during certain tasks, and thus extracts high-level characteristics of users’ responses while the users are instructed to freely look at videos, photos of faces, or other specific types of stimuli. Prior work includes analysis of scan paths and arch densities [7], areas of interest on human faces [14], graph matching [35] and fixation density maps [36].

As summarized in Table 2, existing work in this category mostly achieves Equal Error Rates higher than 15%, which is likely due to complex features being more dependent on varying cognitive and physiological states of the user. Furthermore, in order to acquire sufficient data to extract complex features, these systems require often long authentication times (measured in tens of seconds!), so further improvements are needed before they can be applied to real-world systems.

Low-level Characteristics. On the other hand, motivated by psychological and neurophysiological research [8] that suggests stable differences between users [44], several authors researched systems that use low-level characteristics of users’ eye movements as features for discrimination, such as eye movement velocity profiles, sizes of fixation areas, saccade latencies, etc.

Kasprowski is one of the first authors to start systematically researching the low-level characteristics of user’s gaze for authentication. In his initial paper [22] and corresponding PhD thesis [19], he proposes using features such as the distance between the left and right eye-gaze, Fourier and wavelet transforms of the raw gaze signal, and average velocity directions. The used stimulus consists of 9 LED lights arranged in a 3x3 grid, where the position of the single active light changes according to a fixed, equally timed sequence, regardless of the user’s gaze. An experimental study showed half total error rates of close to 12%, but with relatively high false reject rates of 22%. In relation to our proposal, such stimulus also leads to eliciting some reflexive saccades, but as Table 2 shows, it results in longer authentication times and higher error rates. This is likely due to periods of time where the user has already gazed at the light, but is still waiting for the position of the active LED to change. Finally, the authors propose, organize and describe two yearly competitions in eye movements verification and identification using their datasets [20, 21], which have further increased the research interest in gaze-based authentication.

Komogortsev proposes modeling the physiological properties of individuals’ oculomotor plant [27] during multiple

horizontal saccades and using the estimated model parameters as features for classification. Related work by Holland et al. [17] provides an insight into performances of multiple features such as fixation counts and durations during text reading and combines these two approaches to achieve an EER of 23%, while the newer research [16] provides an additional analysis of 13 classification features based on fixations and saccades and achieves an EER of 16.5%.

In contrast to point-of-entry authentication, Eberz et al. [13] propose using 21 low-level characteristics of eye movements to continuously re-authenticate users, regardless of their current task. For one parameter combination, the authors achieve Equal Error Rates of 7.8% when 40 seconds are chosen as a period before making the first decision. Primarily because of requirement of task independence in a continuous authentication scenario, potential replay attacks remain a serious vulnerability. If the attacker is able to capture even a very short recording of legitimate user's gaze, he can continuously rewind and replay it back to the gaze tracking device, and this causes the system to (correctly!) accept the received eye movements as coming from a legitimate user.

9. DISCUSSION

Advanced Attacks. A more sophisticated attacker could build a model of a legitimate user's eye movements to successfully respond to a given challenge. However, we argue that performing such attacks is not straightforward and requires a higher level of complexity than simply replaying a biometric sample.

Firstly, the adversary is likely to be solving a harder problem than the system; while the system needs to build a discriminative model that allows making a binary decision about user's identity, the adversary needs to actually generate eye movements which correspond to the legitimate user. An indication of the difficulty of artificially creating eye-movements can be found in work by Komogortsev et al. [28], which evaluated the complexity of a significantly simpler problem: artificially generating 1-dimensional eye movements. The paper showed that those movements could be distinguished from natural recordings with high accuracy; creating realistic 2D eye-movements that correspond to a specific user is likely to be significantly harder.

Secondly, by using a challenge-response type of protocol, we ensure that the potential generative model of legitimate user's eye movements must be able to output results interactively and in real-time since the stimulus is not known in advance. This requires additional level of sophistication that is not needed for replay attacks, since the adversary needs to not only control the gaze tracking channel, but to also observe and analyze the visual channel.

In conclusion, while we acknowledge that an adversary in possession of multiple authentication attempts can attempt different targeted attacks, we argue that they require significantly higher level of sophistication and dedication than what is needed to simply replay an acquired biometric sample. Therefore, we believe that successfully preventing replay attacks, as the most applicable threat vector against biometrics, is an important step towards their widespread deployment.

How fast is fast enough? The most common form of user authentication in today's world is password-based. Passwords are relatively fast to type and easy to implement. Considering their prevalence and simplicity, we use passwords as an informal benchmark in terms of authentication times and input error rates to assess the future potential of gaze-based authentication based on reflexive eye behavior.

Over the last few years, a wide range of studies on password authentication have been published, several of those focusing on evaluation of entry times for different password generation strategies, as well as input and recall error rates. From a usability standpoint, we believe that a recent paper by Shay et al. [39] provides an estimate of password usage in a realistic setting and with a large number of users. The authors evaluate multiple password-composition policies by running an online experiment with 8,143 participants, who are asked to create, remember and recall different passwords. Depending on the required password complexity, the median input times varied from 11.6 to 16.2 seconds, while input error rates ranged between 4% and 7%. The authors also note that more than 20% of participants had problems recalling their password and more than 35% of participants stated that remembering a password was hard.

Considering these findings, we believe that our results, namely a median authentication time of 5 seconds and an equal error rate of 6.3%, are highly comparable to input times and error rates on passwords: on average, a successful authentication attempt with our proposed system does not take longer than typing a password, with an added benefit that users need neither learn nor recall any information or procedure.

10. CONCLUSION

Building upon the core idea of using reflexive human behavior for authentication, in this paper we designed an interactive visual stimulus for rapidly eliciting standardized reflexive eye movements, and showed how such stimulus can be used to construct a fast challenge-response biometric system. Based on a series of user experiments, we showed that our stimulus indeed elicits predominately reflexive saccades, which are automatic responses that only pose low cognitive load on the user. As a result of using reflexive behavior that is fast and stable, we show that our authentication system achieves fast authentication times (median of 5 seconds) and low error rates (6.3% EER).

Most importantly, however, our proposed authentication method shows resilience against replay attacks, a property difficult to achieve with most biometrics. Evaluation shows that the system is able to detect the replay of recorded eye traces with very high probability of 99.94%, thus preventing one of the most applicable attacks on biometric systems.

Considering the recent proliferation of reliable and affordable eye tracking devices, we believe that achieving fast and reliable gaze-based authentication is of broad interest and we consider our work to be an important step in this direction.

Finally, this paper opens several interesting questions for future work, such as could reflexive human behavior be exploited in other biometric modalities, or how could reflexive behavior of human visual system be used to support other authentication methods?

Acknowledgements. We thank the anonymous reviewers for their valuable feedback. Ivo Sluganovic is supported by the UK EPSRC doctoral studentship, Scatcherd European Scholarship, and the Frankopan Fund. Authors wish to thank Armasuisse for support with gaze tracking equipment.

11. REFERENCES

- [1] W. W. Abbott and A. A. Faisal. Ultra-low-cost 3D gaze estimation: an intuitive high information throughput compliment to direct brain-machine interfaces. *Journal of Neural Engineering*, 9(4), 2012.
- [2] R. Abrams, D. E. Meyer, and S. Kornblum. Speed and accuracy of saccadic eye movements: characteristics of impulse variability in the oculomotor system. *Journal of experimental psychology. Human perception and performance*, 15(3), 1989.
- [3] T. Bahill, M. R. Clark, and L. Stark. The main sequence, a tool for studying human eye movements. *Mathematical Biosciences*, 24(3-4):191–204, 1975.
- [4] T. Bahill and T. Laritz. Why Can't Batters Keep Their Eyes on the Ball? *American Scientist*, (May - June), 1984.
- [5] A. Boehm, D. Chen, M. Frank, L. Huang, C. Kuo, T. Lolic, I. Martinovic, and D. Song. Safe: Secure authentication with face and eyes. In *Privacy and Security in Mobile Systems (PRISMS), 2013 International Conference on*, June 2013.
- [6] A. Bulling, F. Alt, and A. Schmidt. Increasing the security of gaze-based cued-recall graphical passwords using saliency masks. In *CHI*, 2012.
- [7] V. Cantoni, C. Galdi, M. Nappi, M. Porta, and D. Riccio. Gant: Gaze analysis technique for human identification. *Pattern Recognition*, 48(4), 2015.
- [8] M. S. Castelhana and J. M. Henderson. Stable individual differences across images in human saccadic eye movements. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 62(1):1–14, 2008.
- [9] J. E. S. Choi, P. a. Vaswani, and R. Shadmehr. Vigor of movements and the cost of time in decision making. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 34(4), 2014.
- [10] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20:273–297, 1995.
- [11] A. De Luca, M. Denzel, and H. Hussmann. Look into my eyes! Can you guess my password? In *Proceedings of the 5th Symposium on Usable Privacy and Security*, SOUPS '09, New York, NY, USA, 2009. ACM.
- [12] F. Di Russo, S. Pitzalis, and D. Spinelli. Fixation stability and saccadic latency in elite shooters. *Vision Research*, 43(17), 2003.
- [13] S. Eberz, K. B. Rasmussen, V. Lenders, and I. Martinovic. Preventing Lunchtime Attacks: Fighting Insider Threats With Eye Movement Biometrics. In *Proceedings of the 2015 Networked and Distributed System Security Symposium.*, 2015.
- [14] C. Galdi, M. Nappi, D. Riccio, V. Cantoni, and M. Porta. A new gaze analysis based soft-biometric. *Lecture Notes in Computer Science*, 7914 LNCS, 2013.
- [15] L. R. Gottlob, M. T. Fillmore, and B. D. Abrams. Age-group differences in saccadic interference. *The journals of gerontology. Series B, Psychological sciences and social sciences*, 62(2):85–89, 2007.
- [16] C. Holland and O. Komogortsev. Complex eye movement pattern biometrics: Analyzing fixations and saccades. In *Biometrics (ICB), 2013 International Conference on*, June 2013.
- [17] C. Holland and O. V. Komogortsev. Biometric identification via eye movement scanpaths in reading. *2011 International Joint Conference on Biometrics, IJCB 2011*, 2011.
- [18] K. Holmqvist, M. Nyström, R. Andersson, R. Dewhurst, J. Halszka, and J. van de Weijer. *Eye Tracking : A Comprehensive Guide to Methods and Measures*. Oxford University Press, 2011.
- [19] P. Kasprowski. Human Identification Using Eye Movements. *Institute of Computer Science*, 2004.
- [20] P. Kasprowski. The Second Eye Movements Verification and Identification Competition. In *IEEE & IAPR International Joint Conference on Biometrics*, 2014.
- [21] P. Kasprowski, O. V. Komogortsev, and A. Karpov. First eye movement verification and identification competition at BTAS 2012. *2012 IEEE 5th International Conference on Biometrics: Theory, Applications and Systems, BTAS 2012*, (Btas), 2012.
- [22] P. Kasprowski and J. Ober. Eye Movements in Biometrics. *Biometrics*, 3087 / 200, 2003.
- [23] Katharine Byrne. MSI & Tobii join forces to create eye-tracking gaming laptop, 2015.
- [24] A. Klin, W. Jones, R. Schultz, F. Volkmar, and D. Cohen. Visual fixation patterns during viewing of naturalistic social situations as predictors of social competence in individuals with autism. *Archives of general psychiatry*, 59:809–816, 2002.
- [25] T. Kocejko and J. Wtorek. *Information Technologies in Biomedicine: Third International Conference, ITIB 2012, Gliwice, Poland, June 11-13, 2012. Proceedings*, chapter Gaze Pattern Lock for Elders and Disabled, pages 589–602. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- [26] O. V. Kolesnikova, L. V. Tereshchenko, a. V. Latanov, and V. V. Shulgovskii. *Neuroscience and Behavioral Physiology*, 40(8):869–876, 2010.
- [27] O. V. Komogortsev, U. K. S. Jayarathna, C. R. Aragon, and M. Mechehou. Biometric Identification via an Oculomotor Plant Mathematical Model. *Eye Tracking Research & Applications (ETRA) Symposium*, 2010.
- [28] O. V. Komogortsev, A. Karpov, and C. D. Holland. Attack of Mechanical Replicas : Liveness Detection With Eye Movements. *IEEE TIFS*, 10(4), 2015.
- [29] M. Kumar, T. Garfinkel, D. Boneh, and T. Winograd. Reducing shoulder-surfing by using gaze-based password entry. In *Proceedings of the 3rd Symposium on Usable Privacy and Security*, SOUPS '07, New York, NY, USA, 2007. ACM.
- [30] M. F. Land. Oculomotor behaviour in vertebrates and invertebrates. *The Oxford handbook of eye movements*, 1, 2011.
- [31] I. E. Lazarev and a. V. Kirenskaia. Effect of eye dominance on saccade characteristics and slow EEG potentials. *Fiziologiya cheloveka*, 34(2):23–33, 2008.
- [32] E. Miluzzo, T. Wang, A. T. Campbell, and a. C. M. S. I. G. o. D. Communication. EyePhone: Activating Mobile Phones with Your Eyes. *Workshop on Networking, Systems, and Applications on Mobile Handhelds (MobiHeld)*, 2010.
- [33] M. Nystrom and K. Holmqvist. An adaptive algorithm for fixation, saccade, and glissade detection in eyetracking data. *Behavior Research Methods*, 42(1), 2010.
- [34] T. Poitschke, F. Laquai, S. Stamboliev, and G. Rigoll. Gaze-based interaction on multiple displays in an automotive environment. In *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, 2011.
- [35] I. Rigas, G. Economou, and S. Fotopoulos. Biometric identification based on the eye movements and graph matching techniques. *Pattern Recognition Letters*, 33(6), 2012.
- [36] I. Rigas and O. V. Komogortsev. Biometric Recognition via Probabilistic Spatial Projection of Eye Movement Trajectories in Dynamic Visual Environments. *IEEE TIFS*, 9(10), 2014.
- [37] U. Saeed. Eye movements during scene understanding for biometric identification. *Pattern Recognition Letters*, 59, 2015.
- [38] SensoMotoric Instruments GmbH. SMI RED500 Technical Specification. Technical report, SensoMotoric Instruments GmbH, Teltow, Germany, 2011.
- [39] R. Shay, L. F. Cranor, S. Komanduri, A. L. Durity, P. S. Huh, M. L. Mazurek, S. M. Segreti, B. Ur, L. Bauer, and N. Christin. Can long passwords be secure and usable? *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, 2014.
- [40] P. Sumner. Determinants of saccade latency. In *Oxford handbook of eye movements*, volume 22, pages 411–424. 2011.
- [41] R. Walker, D. G. Walker, M. Husain, and C. Kennard. Control of voluntary and reflexive saccades. *Experimental Brain Research*, 130(4):540–544, Feb. 2000.
- [42] J. Weaver, K. Mock, and B. Hoanca. Gaze-based password authentication through automatic clustering of gaze points. *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, 2011.
- [43] Y. Zhang, Z. Chi, and D. Feng. An Analysis of Eye Movement Based Authentication Systems. *International Conference on Mechanical Engineering and Technology (ICMET-London 2011)*, 2011.
- [44] M. Zhang, Y. Laurikkala, J. Juhola. Biometric verification of a subject with eye movements, with special reference to temporal variability in saccades between a subject's measurements. *Int. J. Biometrics*, 6(1), 2014.