



Privacy Preserving Gaze Estimation using Synthetic Images via a Randomized Encoding Based Framework

Efe Bozkir*
efe.bozkir@uni-tuebingen.de
Human-Computer Interaction
University of Tübingen, Germany

Ali Burak Ünal*
uenal@informatik.uni-tuebingen.de
Methods in Medical Informatics
University of Tübingen, Germany

Mete Akgün
mete.akguen@uni-tuebingen.de
Methods in Medical Informatics,
Translational Bioinformatics
University of Tübingen, Germany

Enkelejda Kasneci
enkelejda.kasneci@uni-tuebingen.de
Human-Computer Interaction
University of Tübingen, Germany

Nico Pfeifer†
pfeifer@informatik.uni-tuebingen.de
Methods in Medical Informatics
University of Tübingen, Germany

ABSTRACT

Eye tracking is handled as one of the key technologies for applications that assess and evaluate human attention, behavior, and biometrics, especially using gaze, pupillary, and blink behaviors. One of the challenges with regard to the social acceptance of eye tracking technology is however the preserving of sensitive and personal information. To tackle this challenge, we employ a privacy-preserving framework based on randomized encoding to train a Support Vector Regression model using synthetic eye images privately to estimate the human gaze. During the computation, none of the parties learn about the data or the result that any other party has. Furthermore, the party that trains the model cannot reconstruct pupil, blinks or visual scanpath. The experimental results show that our privacy-preserving framework is capable of working in real-time, with the same accuracy as compared to non-private version and could be extended to other eye tracking related problems.

CCS CONCEPTS

• **Security and privacy**; • **Computing methodologies** → *Kernel methods*; • **Human-centered computing**;

KEYWORDS

privacy preserving machine learning, gaze estimation, randomized encoding, eye tracking, human computer interaction

ACM Reference Format:

Efe Bozkir, Ali Burak Ünal, Mete Akgün, Enkelejda Kasneci, and Nico Pfeifer. 2020. Privacy Preserving Gaze Estimation using Synthetic Images via a

Randomized Encoding Based Framework. In *Symposium on Eye Tracking Research and Applications (ETRA '20 Short Papers)*, June 2–5, 2020, Stuttgart, Germany. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3379156.3391364>

1 INTRODUCTION

Recent advances in the fields of Head-Mounted-Display (HMD) technology, computer graphics, augmented reality (AR), and eye tracking enabled numerous novel applications. One of the most natural and non-intrusive ways of interaction with HMDs or smart glasses is achieved by gaze-aware interfaces using eye tracking. However, it is possible to derive a lot of sensitive and personal information from eye tracking data such as intentions, behaviors, or fatigue since eyes are not fully controlled in a conscious way.

It has been shown that cognitive load [Appel et al. 2018; Chen and Epps 2014], visual attention [Bozkir et al. 2019], stress [Kübler et al. 2014], task identification [Borji and Itti 2014], skill level assessment and expertise [Castner et al. 2018; Eivazi et al. 2017; Liu et al. 2009], human activities [Braunagel et al. 2017; Steil and Bulling 2015], biometric information and authentication [Abdrabou et al. 2019; Kinnunen et al. 2010; Komogortsev and Holland 2013; Komogortsev et al. 2010; Zhang et al. 2018], or personality traits [Berkovsky et al. 2019] can be obtained using eye tracking data. Since highly sensitive information can be derived from eye tracking data, it is not surprising that HMDs or smart glasses have not been adopted by large communities yet. According to a recent survey [Steil et al. 2019a], people agree to share their eye tracking data only when it is co-owned by a governmental health-agency or is used for research purposes. This indicates that people are hesitant about sharing their eye tracking data in commercial applications. Therefore, there is a likelihood that larger communities could adopt HMDs or smart glasses if privacy-preserving techniques are applied in the eye tracking applications. The reasons why privacy preserving schemes are needed for eye tracking are discussed in [Liebling and Preibusch 2014] extensively. However, until now, there are not many studies in privacy-preserving eye tracking. Recently, a method to detect privacy sensitive everyday situations [Steil et al. 2019b], an approach to degrade iris authentication while keeping the gaze tracking utility in an acceptable accuracy [John et al. 2019], and differential privacy based techniques to protect personal information on heatmaps and eye movements [Liu et al. 2019; Steil et al.

*Both authors contributed equally to this research.

†Also affiliated with Statistical Learning in Computational Biology, Max Planck Institute for Informatics, Saarbrücken, Germany

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ETRA '20 Short Papers, June 2–5, 2020, Stuttgart, Germany

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7134-6/20/06...\$15.00

<https://doi.org/10.1145/3379156.3391364>

2019a] are introduced. While differential privacy can be applied to eye tracking data for various tasks, it introduces additional noise on the data which causes decrease in the utility [Liu et al. 2019; Steil et al. 2019a], and it might lead to less accurate results in computer vision tasks, such as gaze estimation or activity recognition.

In light of the above, function-specific privacy models are required. In this work, we focus on the gaze estimation problem as a proof-of-concept by using synthetic data including eye landmarks and ground truth gaze vectors. However, the same privacy-preserving approach can be extended to any feature-based, eye tracking problem such as intention, fatigue, or activity detection, in HMD or unconstrained setups due to the demonstrated real-time working capabilities. In our study, the gaze estimation task is solved by using Support Vector Regression (SVR) models in a privacy-preserving manner by computing the dot product of eye landmark vectors to obtain the kernel matrix of the SVR for a scenario, where two parties have the eye landmark data, each of which we call *input-party*, and one *function-party* that trains a prediction model on the data of the input-parties. This scenario is relevant when the input-parties use eye tracking data to improve the accuracy of their models and do not share the data due to the privacy concerns. To this end, we utilize a framework employing randomized encoding [Ünal et al. 2019]. In the computation, neither the eye images nor the extracted features are revealed to the function-party directly. Furthermore, the input-parties do not infer the raw eye tracking data or result of the computation. Eye images that are used for training and testing are rendered using UnityEyes [Wood et al. 2016] synthetically and 36 landmark-based features [Park et al. 2018] are used. To the best of our knowledge, this is the first work that applies a privacy-preserving scheme based on function-specific privacy models on an eye tracking problem.

2 THREAT MODEL

We assume that the input-parties are semi-honest (honest but curious) that are not allowed to deviate from the protocol description while they try to infer some valuable information about other parties' private inputs using their views of the protocol execution. We also assume that the function-party is malicious and the input-parties and the function-party do not collude.

3 METHODOLOGY

In this section, we discuss the data generation, randomized encoding, and privacy-preserving gaze estimation framework.

3.1 Data Generation

To train and evaluate the gaze estimator, we generate eye images and gaze vectors. As our work is a proof-of-concept and requires high amount of data, synthetic images from UnityEyes [Wood et al. 2016], which is based on the Unity3D, are used. *Camera parameters* and *Eye parameters* are chosen as $(0, 0, 0, 0)$ (fixed camera) and $(0, 0, 30, 30)$ (eyeball pose range parameters in degrees), respectively. 20,000 images are rendered in *Fantastic* quality setting and 512×384 screen resolution. Then, processing and normalization pipeline from [Park et al. 2018] is employed. In the end, we obtain 128×96 sized eye images, 18 eye landmarks including eight iris edge, eight eyelid, one iris center, and one iris-center-eyeball-center vector




Figure 1: Eye landmarks and gaze on a synthetic image.

normalized according to Euclidean distance between eye corners, and gaze vectors using pitch and yaw angles. Final feature vectors consist of 36 elements. Figure 1 shows an example illustration.

3.2 Randomized Encoding

The utilized framework employs randomized encoding (RE) [Applebaum et al. 2006a,b] to compute the dot product of the landmark vectors. The dot product is needed to compute kernel matrix of the SVR which is later used for training the gaze estimator and validation of the framework.

In the randomized encoding, the computation of a function $f(x)$ is performed by a randomized function $\hat{f}(x; r)$ where x is the input value, which corresponds to eye landmarks in our setup, and r is the random value. The idea is to encode the original function by using random value(s) such that the combination of the components of the encoding reveals only the output of the original function. In the framework, the computation of the dot product is accomplished by utilizing the decomposable and affine randomized encoding (DARE) of addition and multiplication [Applebaum 2017]. The encoding of multiplication is as follows.

Definition 1 (Perfect RE for Multiplication [Applebaum 2017]). A multiplication function is defined as $f_m(x_1, x_2) = x_1 \cdot x_2$ over a ring R . One can perfectly encode the f_m by employing the DARE $\hat{f}_m(x_1, x_2; r_1, r_2, r_3)$:

$$\hat{f}_m(x_1, x_2; r_1, r_2, r_3) = (x_1 + r_1, x_2 + r_2, r_2x_1 + r_3, r_1x_2 + r_1r_2 - r_3),$$

where r_1, r_2 and r_3 are uniformly chosen random values. The recovery of $f_m(x_1, x_2)$ can be accomplished by computing $c_1 \cdot c_2 - c_3 - c_4$ where $c_1 = x_1 + r_1$, $c_2 = x_2 + r_2$, $c_3 = r_2x_1 + r_3$ and $c_4 = r_1x_2 + r_1r_2 - r_3$. The simulation of \hat{f}_m can be done perfectly by the simulator $\text{Sim}(y; a_1, a_2, a_3) := (a_1, a_2, a_3, a_1a_2 - y - a_3)$ where a_1, a_2, a_3 are random values.

3.3 Framework

To perform the private gaze estimation task in our scenario, we inspire from the framework as in [Ünal et al. 2019] due to its efficiency compared to other approaches in the literature. The framework is proposed to compute the addition or multiplication of the input values of two input-parties in the function-party by utilizing randomized encoding. We utilize the multiplication operation over the eye landmark vectors to compute the dot product of these vectors to obtain kernel matrix of the SVR in a privacy-preserving way.

We have two input-parties as Alice and Bob, having the eye landmark data as $X \in \mathbb{R}^{n_f \times n_a}$ and $Y \in \mathbb{R}^{n_f \times n_b}$ where n_a and n_b




Figure 2: Overall protocol execution.

represent the number of samples in Alice and Bob, respectively, and n_f is the number of features. In addition to the input-parties, there exists a server that trains a model on the data of the input-parties. $A_{.j}$ for any matrix A represents the j -th column of the corresponding matrix and “ \odot ” represents the element-wise multiplication of the vectors. As a first step, Alice creates a uniformly chosen random value $r_3 \in \mathbb{R}$ and two vectors $r_1, r_2 \in \mathbb{R}^{n_f}$ with uniformly chosen random values, which are used to encode the element-wise multiplication of the vectors and shares them with Bob. Afterwards, Bob computes $C_{.j}^2 = Y_{.j} + r_2$ and $C_j^4 = \sum_{d=1}^{n_f} (r_1 \odot Y_{.j} + r_1 \odot r_2)_d - r_3$, $\forall j \in \{1, \dots, n_b\}$ where $C^2 \in \mathbb{R}^{n_f \times n_b}$ and $C^4 \in \mathbb{R}^{n_b}$. Meanwhile, Alice computes $C_{.i}^1 = X_{.i} + r_1$ and $C_i^3 = \sum_{d=1}^{n_f} (r_2 \odot X_{.i})_d + r_3$, $\forall i \in \{1, \dots, n_a\}$ where $C^1 \in \mathbb{R}^{n_f \times n_a}$ and $C^3 \in \mathbb{R}^{n_a}$. Input-parties send their share of the encoding to the server with the gram matrix of their samples, which is the dot product among their samples. Then, the server computes the dot product between samples of Alice and Bob to complete the missing part of the gram matrix of all samples. To achieve this, the server computes $k_{ij} = \sum_{d=1}^{n_f} (C_{.i}^1 \odot C_{.j}^2)_d - C_i^3 - C_j^4$, $\forall i \in \{1, \dots, n_a\}$ and $\forall j \in \{1, \dots, n_b\}$ where k_{ij} is the i -th row j -th column entry of the gram matrix between the samples of the input-parties. Once the server has all components of the gram matrix, it constructs the complete gram matrix K by simply concatenating the parts of it. In our solution, Alice and Bob send to the server (C^1, C^3) and (C^2, C^4) tuples, respectively. These components reveal nothing but only the gram matrix of the samples after decoding. Furthermore, the input-parties shuffle their raw data before the computation to avoid the possibility of private information leakage such as the behavior of the person due to the nature of the visual sequence information. The overall flow is summarized in Figure 2.

After having the complete gram matrix for all samples that Alice and Bob have, the server uses it as a kernel matrix as if it was computed by the linear kernel function on pooled data. Additionally, it is also possible to compute a kernel matrix as if it was computed by the polynomial or radial basis kernel function (RBF) by utilizing

the resulting gram matrix. As an example, the calculation of RBF from the gram matrix is as follows.

$$K(x, y) = \exp \left(- \frac{\|x \cdot x - 2x \cdot y + y \cdot y\|^2}{2\sigma^2} \right),$$

where “ \cdot ” represents the dot product of vectors, which is possible to obtain from the gram matrix, and σ is the parameter utilized to adjust the similarity level. Once the desired kernel matrix is computed, it is possible to train an SVR model by employing the computed kernel matrix to estimate the gaze. In the process of the computation of the dot product, the amount of data transferred among parties is $(n_f n_a + n_f n_b + n_a + n_b + 2n_f) \times d$ bytes where d is the size of one data unit.

4 SECURITY ANALYSIS

A semi-honest adversary who corrupts any of the input-parties cannot learn anything about the private inputs of the other input-party. During the protocol execution, two vectors of random values and a single random value are sent from Alice to Bob. The views of the input-parties consist only of vectors with random values. Using these random values, it is not possible for one party to infer something about the other party’s private inputs [Unal et al. 2019].

Theorem 1. A malicious adversary \mathcal{A} corrupting the function-party learns nothing more than the result of gram matrix. It is computationally infeasible for \mathcal{A} to infer any information about the input-parties’ data X and Y as long as Perfect RE multiplication is semantically secure (Definition 1).

PROOF. We first show the correctness of our solution. We assume $n_f = 2$ and encode the function $f_d(x, y) = x_1 y_1 + x_2 y_2$ over some finite ring R by the following DARE:

$$\begin{aligned} \hat{f}_d(x, y; r) &= (x_1 + r_{11}, y_1 + r_{12}, x_2 + r_{21}, y_2 + r_{22}, \\ &\quad r_{12}x_1 + r_{22}x_2 + r_3, \\ &\quad r_{11}y_1 + r_{11}r_{12} + r_{21}y_2 + r_{21}r_{22} - r_3) \end{aligned}$$

Given an encoding $(c_1, c_2, c_3, c_4, c_5, c_6)$, $f_d(x, y)$ is recovered by computing $c_1 c_2 + c_2 c_4 + c_5 + c_6$.

By the concatenation lemma in [Applebaum 2017], we can divide c_5 and c_6 into n_f shares by using n_f random values instead of a single r_3 value.

$$\begin{aligned} \hat{f}_d(x, y; r) &= (x_1 + r_{11}, y_1 + r_{12}, r_{12}x_1 + r_{13}, r_{11}y_1 + r_{11}r_{12} - r_{13}, \\ &\quad x_2 + r_{21}, y_2 + r_{22}, r_{22}x_2 + r_{23}, r_{21}y_2 + r_{21}r_{22} - r_{23}) \end{aligned}$$

Given an encoding $(c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_8)$,

$$\begin{aligned} \hat{f}_m(x_1, y_1; r) &= (c_1, c_2, c_3, c_4) \\ \hat{f}_m(x_2, y_2; r) &= (c_5, c_6, c_7, c_8) \end{aligned}$$

By the concatenation lemma in [Applebaum 2017], $\hat{f}_d(x, y; r) = (\hat{f}_m(x_1, y_1; r), \hat{f}_m(x_2, y_2; r))$ perfectly encodes the function $f_d(x, y)$ if Perfect RE multiplication is semantically secure.

After showing the correctness, we analyze the security with the simulation paradigm. In the simulation paradigm, there is a simulator who generates the view of a party in the execution. A party’s input and output must be given to the simulator to generate

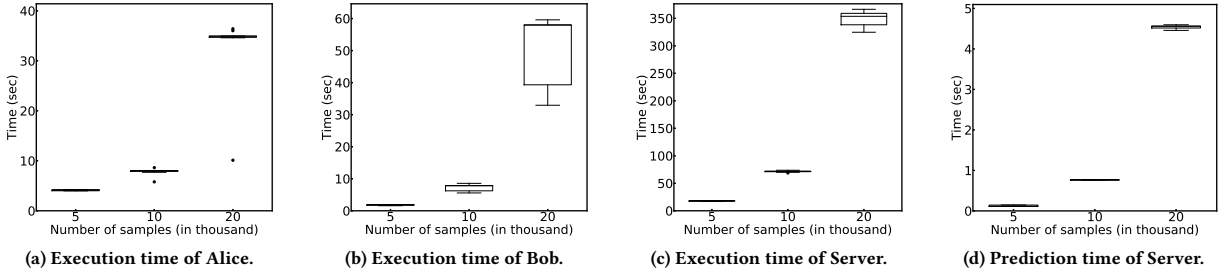


Figure 3: The execution time of (a) Alice, (b) Bob and (c) the server are given. We also demonstrate (d) the time required for the prediction of the test samples, which are 20% of the total number of samples in each case.

its view. Thus, security is formalized by saying that a party's view can be simulatable given its input and output and the parties learn nothing more than what they can derive from their input and prescribed output.

The function-party \mathcal{F} does not have any input and output. A simulator \mathcal{S} can generate the views of incoming messages received by \mathcal{F} . \mathcal{S} creates four vectors C^1, C^2, C^3, C^4 with uniformly distributed random values using a pseudorandom number generator G' . Finally, \mathcal{S} outputs $\{C^1, C^2, C^3, C^4\}$.

In the execution of the protocol π , \mathcal{A} receives four messages which are masked with uniformly random values generated using a pseudorandom number generator G . The view of \mathcal{A} includes $\{C^1, C^2, C^3, C^4\}$. The distribution over G is statistically close to the distribution over G' . This implies that

$$\{S(C^1, C^2, C^3, C^4)\} \stackrel{c}{\equiv} \{view_{\pi}^{\mathcal{A}}(C^1, C^2, C^3, C^4)\}$$

□

5 RESULTS

To demonstrate the performance, we conduct experiments on a PC equipped with Intel Core i7-7500U with 2.70 GHz processor and 16 GB memory RAM. We employ varying sizes of eye landmark data, that are 5,000, 10,000 and 20,000 samples of which one-fifth is the test data and we split the data between the input-parties equally. The framework allows us to optimize the parameters of the model in the server without further communicating with the input-parties. Thanks to this, we utilize 5-fold cross-validation to optimize the parameters, which are the similarity adjustment parameter $\gamma \in \{2^{-3}, 2^{-2}, \dots, 2^4\}$ of the Gaussian RBF kernel, the misclassification penalty parameter $C \in \{2^{-3}, 2^{-2}, \dots, 2^3\}$, and the tolerance parameter $\epsilon \in \{0.005, 0.01, 0.05, 0.1, 0.5, 1\}$ of SVR. After parameter optimization, we repeat the experiment on varying sizes of eye landmark data with the optimal parameter set 10 times to assess the execution time. To evaluate the gaze estimation results, we employ mean angular error in the same way as in [Park et al. 2018]. Table 1 demonstrates the relationship between the dataset size and the resulting mean angular error. Since no additional noise is introduced during the computation of the kernel matrix, the results from our privacy-preserving framework are the same with the non-private ones. The mean angular errors are lower compared to the state-of-the-art gaze estimation techniques since we use synthetic data and fixed camera position during image rendering.

Table 1: The mean angular errors for varying dataset sizes.

# of samples	Mean angular error
5k	0.21
10k	0.18
20k	0.17

The amount of time to train and test the models increases as the sample sizes increase since computation requirements get larger. The increment in the dataset size increases the communication cost among parties. The execution times of all parties for 10 runs with the optimal parameters are shown in Figure 3. We also demonstrate the amount of time to predict the test samples, which corresponds to one-fifth of the total number of samples to emphasize the real-time working capabilities. In the experiment with 20,000 samples, for instance, we spend ≈ 4.5 seconds to predict 4,000 test samples, which corresponds to 1.125 ms per sample. When the current sampling frequencies of eye trackers are taken into consideration, it is possible to deploy and use the framework to estimate gaze if an optimized communication between the parties is established.

6 CONCLUSION

In this work, we utilized a framework based on randomized encoding to estimate human gaze in a privacy-preserving way and in real-time. Our solution can provide improved gaze estimation if input-parties want to use each other's data for different reasons such as to account for genetic structural differences in the eye region. None of the input-parties has the access to the eye landmark data of the others or the result of the computation in the function party, while the function-party cannot infer anything about the data of the input-parties. Temporal information of the visual scanpath, pupillary, or blinks cannot be reconstructed due to the shuffling of the data, and lack of sensory information and direct access to the eye landmarks. Our solution works in real-time, hence it could be deployed along with HMDs for different use-cases and extended to similar eye tracking related problems if similar amount of features is used. To the best of our knowledge, this is the first work based on function-specific privacy models in the eye tracking domain. The number of parties is a limitation of our solution. Thus, as future work we will extend our work to a larger number of parties.

REFERENCES

- Yasmeen Abdrabou, Mohamed Khamis, Rana Mohamed Eisa, Sherif Ismail, and Amr Elmougy. 2019. Just Gaze and Wave: Exploring the Use of Gaze and Gestures for Shoulder-surfing Resilient Authentication. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications (ETRA '19)*. ACM, New York, NY, USA, Article 29, 10 pages. <https://doi.org/10.1145/3314111.3319837>
- Tobias Appel, Christian Scharinger, Peter Gerjets, and Enkelejd Kasneci. 2018. Cross-subject Workload Classification Using Pupil-related Measures. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications (ETRA '18)*. ACM, New York, NY, USA, Article 4, 8 pages. <https://doi.org/10.1145/3204493.3204531>
- Benny Applebaum. 2017. Garbled circuits as randomized encodings of functions: a primer. In *Tutorials on the Foundations of Cryptography*. Springer, 1–44.
- Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. 2006a. Computationally private randomizing polynomials and their applications. *computational complexity* 15, 2 (2006), 115–162.
- Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. 2006b. Cryptography in NC⁰. *SIAM J. Comput.* 36, 4 (2006), 845–888.
- Shlomo Berkovsky, Ronnie Taib, Irena Koprinska, Eileen Wang, Yucheng Zeng, Jingjie Li, and Sabina Kleitman. 2019. Detecting Personality Traits Using Eye-Tracking Data. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, Article 221, 12 pages. <https://doi.org/10.1145/3290605.3300451>
- Ali Borji and Laurent Itti. 2014. Defending Yarbus: Eye movements reveal observers' task. *Journal of vision* 14 (03 2014). <https://doi.org/10.1167/14.3.29>
- Efe Bozkir, David Geisler, and Enkelejd Kasneci. 2019. Assessment of Driver Attention During a Safety Critical Situation in VR to Generate VR-based Training. In *ACM Symposium on Applied Perception 2019 (SAP '19)*. ACM, New York, NY, USA, Article 23, 5 pages. <https://doi.org/10.1145/3343036.3343138>
- Christian Braunagel, David Geisler, Wolfgang Rosenstiel, and Enkelejd Kasneci. 2017. Online recognition of driver-activity based on visual scanpath classification. *IEEE Intelligent Transportation Systems Magazine* 9, 4 (2017), 23–36.
- Nora Castner, Enkelejd Kasneci, Thomas Kübler, Katharina Scheiter, Juliane Richter, Thérèse Eder, Fabian Hüttig, and Constanze Keutel. 2018. Scanpath Comparison in Medical Image Reading Skills of Dental Students: Distinguishing Stages of Expertise Development. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications (ETRA '18)*. ACM, New York, NY, USA, Article 39, 9 pages. <https://doi.org/10.1145/3204493.3204550>
- Siyuan Chen and Julien Epps. 2014. Using Task-Induced Pupil Diameter and Blink Rate to Infer Cognitive Load. *Human-Computer Interaction* 29 (2014), 390–413.
- Shahram Eivazi, Ahmad Hafez, Wolfgang Fuhl, Hoorieh Afkari, Enkelejd Kasneci, Martin Lehecka, and Roman Bednarik. 2017. Optimal eye movement strategies: a comparison of neurosurgeons gaze patterns when using a surgical microscope. *Acta neurochirurgica* 159, 6 (2017), 959–966.
- Brendan John, Sanjeev Koppal, and Eakta Jain. 2019. EyeVEIL: Degrading Iris Authentication in Eye Tracking Headsets. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications (ETRA '19)*. ACM, New York, NY, USA, Article 37, 5 pages. <https://doi.org/10.1145/3314111.3319816>
- Tomi Kinnunen, Filip Sedlak, and Roman Bednarik. 2010. Towards Task-independent Person Authentication Using Eye Movement Signals. In *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications (ETRA '10)*. ACM, New York, NY, USA, 187–190. <https://doi.org/10.1145/1743666.1743712>
- Oleg V. Komogortsev and Corey D. Holland. 2013. Biometric authentication via complex oculomotor behavior. In *2013 IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*. IEEE, 1–8. <https://doi.org/10.1109/BTAS.2013.6712725>
- Oleg V. Komogortsev, Sampath Jayarathna, Cecilia R. Aragon, and Mechehouli Mahmoud. 2010. Biometric Identification via an Oculomotor Plant Mathematical Model. In *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications (ETRA '10)*. ACM, New York, NY, USA, 57–60. <https://doi.org/10.1145/1743666.1743679>
- Thomas C. Kübler, Enkelejd Kasneci, Wolfgang Rosenstiel, Ulrich Schiefer, Katja Nagel, and Elena Papageorgiou. 2014. Stress-indicators and exploratory gaze for the analysis of hazard perception in patients with visual field loss. *Transportation Research Part F: Traffic Psychology and Behaviour* 24 (2014), 231–243.
- Daniel J. Liebling and Sören Preibusch. 2014. Privacy Considerations for a Pervasive Eye Tracking World. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication (UbiComp '14 Adjunct)*. ACM, New York, NY, USA, 1169–1177. <https://doi.org/10.1145/2638728.2641688>
- Ao Liu, Lirong Xia, Andrew Duchowski, Reynold Bailey, Kenneth Holmqvist, and Eakta Jain. 2019. Differential Privacy for Eye-tracking Data. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications (ETRA '19)*. ACM, New York, NY, USA, Article 28, 10 pages. <https://doi.org/10.1145/3314111.3319823>
- Yan Liu, Pei-Yun Hsueh, Jennifer Lai, Mirweis Sangin, Marc-Antoine Nussli, and Pierre Dillenbourg. 2009. Who is the expert? Analyzing gaze data to predict expertise level in collaborative applications. In *2009 IEEE International Conference on Multimedia and Expo*. IEEE, 898–901. <https://doi.org/10.1109/ICME.2009.5202640>
- Seonwook Park, Xucong Zhang, Andreas Bulling, and Otmar Hilliges. 2018. Learning to Find Eye Region Landmarks for Remote Gaze Estimation in Unconstrained Settings. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications (ETRA '18)*. ACM, New York, NY, USA, Article 21, 10 pages. <https://doi.org/10.1145/3204493.3204545>
- Julian Steil and Andreas Bulling. 2015. Discovery of Everyday Human Activities from Long-term Visual Behaviour Using Topic Models. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 75–85. <https://doi.org/10.1145/2750858.2807520>
- Julian Steil, Inken Hagedstedt, Michael Xuelin Huang, and Andreas Bulling. 2019a. Privacy-aware Eye Tracking Using Differential Privacy. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications (ETRA '19)*. ACM, New York, NY, USA, Article 27, 9 pages. <https://doi.org/10.1145/3314111.3319915>
- Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2019b. PrivacEye: Privacy-preserving Head-mounted Eye Tracking Using Egocentric Scene Image and Eye Movement Features. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications (ETRA '19)*. ACM, New York, NY, USA, Article 26, 10 pages. <https://doi.org/10.1145/3314111.3319913>
- Ali Burak Ünal, Mete Akgün, and Nico Pfeifer. 2019. A Framework with Randomized Encoding for a Fast Privacy Preserving Calculation of Non-linear Kernels for Machine Learning Applications in Precision Medicine. In *Cryptology and Network Security*. Springer International Publishing, Cham, 493–511.
- Erroll Wood, Tadas Baltrušaitis, Louis-Philippe Morency, Peter Robinson, and Andreas Bulling. 2016. Learning an Appearance-Based Gaze Estimator from One Million Synthesised Images. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications (ETRA '16)*. ACM, New York, NY, USA, 131–138. <https://doi.org/10.1145/2857491.2857492>
- Yongtuo Zhang, Wen Hu, Weitao Xu, Chun Tung Chou, and Jiankun Hu. 2018. Continuous Authentication Using Eye Movement Response of Implicit Visual Stimuli. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 4, Article 177 (Jan. 2018), 22 pages. <https://doi.org/10.1145/3161410>