Differential Privacy for Eye Tracking with Temporal Correlations

EFE BOZKIR*, University of Tübingen, Germany, efe.bozkir@uni-tuebingen.de
ONUR GÜNLÜ*, TU Berlin, Germany, guenlue@tu-berlin.de
WOLFGANG FUHL, University of Tübingen, Germany, wolfgang.fuhl@uni-tuebingen.de
RAFAEL F. SCHAEFER, TU Berlin, Germany, rafael.schaefer@tu-berlin.de
ENKELEJDA KASNECI, University of Tübingen, Germany, enkelejda.kasneci@uni-tuebingen.de

Head mounted displays bring eye tracking into daily use and this raises privacy concerns for users. Privacy-preservation techniques such as differential privacy mechanisms are recently applied to an eye tracking data obtained from such displays; however, standard differential privacy mechanisms are vulnerable to temporal correlations. In this work, a transform coding based differential privacy mechanism is proposed for the first time in the eye tracking literature to further adapt the mechanism to statistics of eye movement features by comparing low-complexity methods. Fourier Perturbation Algorithm, which is a differential privacy mechanism, is extended and a scaling mistake in its proof is corrected. Significant correlation and query sensitivity reductions are illustrated, which provide the best utility-privacy trade-off in the literature for the eye tracking dataset used. Classification accuracies of differentially private eye movements for gender and document-type predictions are evaluated to illustrate that significantly high privacy can be obtained without accuracy loss.

CCS Concepts: • **Security and privacy → Human and societal aspects of security and privacy**; **Privacy-preserving protocols**; • **Human-centered computing → Human computer interaction (HCI)**.

Additional Key Words and Phrases: eye tracking, differential privacy, eye movements, privacy protection, virtual reality.

ACM Reference Format:
Efe Bozkir, Onur Günlü, Wolfgang Fuhl, Rafael F. Schaefer, and Enkelejda Kasneci. 2020. Differential Privacy for Eye Tracking with Temporal Correlations. In *Tallinn '20: Nordic Conference on Human-Computer Interaction, October 25–29, 2020, Tallinn, Estonia.* ACM, New York, NY, USA, 16 pages. https://doi.org/00.0000/0000000.0000000

1 INTRODUCTION

Recent advances in the field of smart glasses, computer hardware, head mounted displays (HMDs), and eye tracking enable easy access to pervasive eye trackers along with modern HMDs. The decrease in the cost of such devices might cause a significant increase in the amount of eye tracking and movement data. Although a large amount of eye tracking data generation is helpful for user assistive and comfort providing tasks especially in the domain of augmented and virtual reality (AR/VR), since eyes are not fully controlled consciously, it is also possible to derive plenty of sensitive information about users from such data.

It has been shown that human activities can be recognized in everyday tasks [31], e.g., during conditional automated driving [5], or while wearing a Google glass [16] with the help of variety of eye tracking features. While pupillometry,

*Both authors contributed equally to this research.

Manuscript submitted to ACM

blink rates, and microsaccades are related to the cognitive load of people [1, 24], mental fatigue can be also detected accurately not only in specific cognitive tasks but also in natural viewing situations by using eye tracking data [35]. Similarly, assessment of situational and visual attention can be made by using eye tracking features [3]. Analysis for expert-novice can be made with the help of eye tracking in areas such as medicine [7] and sports [34]. Furthermore, personality traits and human intent during robotic hand-eye coordination can be also predicted by using eye tracking data [2], [30].

Eye movement data can be also used for biometric authentication, which can be claimed to be a highly sensitive task [13]. A task-independent authentication using eye movement features and Gaussian mixtures is, for example, discussed in [20]. Biometric identification and authentication using an oculomotor plant model and eye movements are introduced in [22, 23]. Moreover, [11] discusses that eye movement features can be used reliably for authentication both in consumer level devices and various real world tasks. Recently, continuous authentication using eye movements for VR headsets is also studied in [36].

One can obtain sensitive information from eye tracking and eye movement data, and it is important to protect it from adversaries. This is also verified by a recent survey conducted in a VR user study [32] that shows that people agree to share their eye tracking data if a governmental health agency is involved in owning data or if the purpose is research. Therefore, privacy-preserving techniques are needed especially on the data sharing side of eye tracking. In the data sharing area, the most natural way to protect privacy of individuals is anonymization by simply removing personal identifiers. However, it is possible to deduce information from anonymized datasets by using other background information and databases via linkage attacks [28]. Thus, more sophisticated techniques for achieving user level privacy are necessary. Differential privacy [9, 10] is one of the most popular and effective frameworks especially in the database applications area. Differential privacy achieves protection of users' privacy by adding randomly generated noise for a given sensitivity and desired privacy. However, high dimensionality of the data and temporal correlations can reduce utility and privacy, respectively. Since eye tracking data and eye movement features are high dimensional, temporally correlated, and usually contain recordings with longer durations as compared to other time series data, it is important to tackle these problems and to provide privacy simultaneously.

In the eye tracking area, both local and global differential privacy can be applied. Local differential privacy adds user level noise to the data but assumes that each user sends its data to a central data collector after adding local noise [8, 12]. For this work, we consider global differential privacy, because there should be a user level data collector and publisher in an eye tracking with HMD scenario.

To apply differential privacy to the eye movement feature data, we evaluate the standard Laplacian Perturbation Algorithm (LPA) [9] and Fourier Perturbation Algorithm (FPA) [29], the latter of which is suitable for time series data such as the eye movement feature signals. The used dataset consists of 52 eye movement features [6] related to fixations, saccades, blinks, and pupil diameter, collected in a VR setup [32]. We propose two different methods that apply the FPA to chunks of data using original eye movement feature signals or consecutive difference signals. While preserving differential privacy using parallel compositions, chunk-based methods decrease query sensitivity and complexity. On the other hand, difference-based method further decreases the correlations between the eye movement features of an individual at different time instances in addition to the decorrelation provided by the FPA algorithm that uses the discrete Fourier transform (DFT) as, e.g. in [14, 15]. The difference-based method provides higher level of privacy since applying differential privacy to correlated data compromises differential privacy and consecutive differences are observed to be less correlated than original consecutive data. Furthermore, we evaluate our methods in the gender and document type classification tasks by using a similar configuration as in [32]. We use the complete data instead of

applying a subsampling step used in [32], applied to reduce the sensitivity and to improve the classification accuracies in the gender and document type. We are the first to propose differential privacy solutions for eye tracking by taking the temporal correlations into account. Furthermore, the previous work in [32] applies the exponential mechanism for differential privacy on the eye movement feature data. The exponential mechanism is useful for situations where the best enumerated response needs to be chosen [10]. However, in the eye movement data, we are not interested in the "best" response but in the value of a feature vector itself. Therefore, we apply the Laplacian mechanism for differential privacy.

Our main contributions are as follows. We propose chunk-based and difference-based differential privacy methods for eye movement features to reduce query sensitivity, complexity, and temporal correlations. Furthermore, we compare our methods with standard techniques such as LPA and FPA by using the inverse of the normalized mean square error (NMSE) as the new utility metric, and gender and document type classification accuracies as the classification metric for the eye movement feature data. We illustrate significantly better performance of our methods as compared to state-of-the-art methods. Our solutions are capable of handling correlation in the data and decrease sensitivity by dividing the data into smaller chunks, so they form the fundamental baseline for privacy-preserving eye tracking using differential privacy.

This paper is organized as follows. In Section 2, we give an overview of the privacy-preserving eye tracking literature. Fundamental definitions and theorems related to the differential privacy and the proposed utility metric are discussed in Section 3. We list the proposed methods in Section 4. Using a public eye tracking dataset, performance of the proposed methods in terms of utility and classification accuracies are shown in Section 5 to significantly improve on all previous methods. Section 6 concludes the paper.

2 RELATED WORK

It is possible to obtain a high amount of sensitive information about individuals by using eye tracking data. However, there are not many works that focus on privacy-preserving eye tracking. In [21], key issues from social, legal, and ethical points of view are addressed to enable socially acceptable body-worn cameras. Eye trackers also lie within the definition of body-worn cameras. More particular to eye tracking research, why privacy considerations are needed for eye tracking world are discussed in [25], where they focus on gaze and pupillometry. Recently, practical solutions are introduced for protecting user identity and sensitive stimuli by degrading iris authentication by introducing optical defocus in the eye tracking setup [17] and automated enable/disable mechanism for eye tracker's first-person camera with the help of mechanical shutter depending on the detection of privacy sensitive content [33], respectively. Furthermore, a function-specific privacy model for privacy-preserving gaze estimation task, which could also be further extended to other eye tracking related problems, is proposed in [4].

While previous methods can be helpful, they are not directly relevant for the data sharing applications. Differential privacy is a useful framework when user identities should be protected. Recently, in the eye tracking community the differential privacy is applied to eye movements in VR [32] and heatmap data [26]. These works do not address the effects of temporal correlations with eye tracking data over time in the differential privacy context. We discuss these effects and propose methods to reduce them.

3 BACKGROUND

In this section, we define the differential privacy, basic mechanisms, and utility metric we use throughout the paper.

Differential privacy is a measure for privacy risk of an individual participating in a database. One classical example is a dataset with weights of N people and a mean function. When an adversary queries the mean function for N people, the average weights of N people is obtained. However, after querying mean function for N people, when adversary queries the dataset for $N-1$ people, the weight of the remaining person will be automatically leaked. Using differential privacy, noise is added to the outcome of a function so that the outcome does not significantly change based on whether or not a random individual participated in the dataset. The amount of noise that is added should be calibrated carefully since high amount of noise might decrease the utility, which means that data could be useless for further analyses when it is differentially private. We define differential privacy in Definition 1.

Definition 1. *ϵ-Differential Privacy (ϵ-DP) [9].* A randomized mechanism M is ϵ-differentially private if for all databases D and D' that differ at most in one element for every $S \subseteq Range(M)$, we have

$$\Pr[M(D) \in S] \leq e^{\epsilon} \Pr[M(D') \in S]. \tag{1}$$

For differential privacy, the variance of the added noise depends on a metric called query sensitivity, so we define the query sensitivity as follows.

Definition 2. *Query sensitivity [9].* For a random query X^n and $w \in \{1, 2\}$, the sensitivity Δ_w of X^n is the smallest number for all databases D and D' that differ at most in one element such that

$$||X^n(D) - X^n(D')||_w \leq \Delta_w(X^n) \tag{2}$$

where

$$||X^n||_w = \sqrt[w]{\sum_{i=1}^{n} \left(|X_i|\right)^w}. \tag{3}$$

We list theorems that are used in the proposed methods.

Theorem 1. Sequential Composition Theorem [27]. Consider n independent mechanisms M_i for $i = 1, 2, ..., n$. If $M_1, M_2, ..., M_n$ are $\epsilon_1, \epsilon_2, ..., \epsilon_n$-differentially private, respectively, then their joint mechanism is $\left(\sum_{i=1}^{n} \epsilon_i\right)$-differentially private.

Theorem 2. Parallel Composition Theorem [27]. Consider n mechanisms as M_i for $i = 1, 2, ..., n$ that are applied to disjoint subsets of a dataset. If $M_1, M_2, ...,$
M_n are $\epsilon_1, \epsilon_2, ..., \epsilon_n$-differentially private, respectively, then their joint mechanism is $\left(\max_{i \in [1,n]} \epsilon_i\right)$-differentially private.

Next, we define the Laplacian Perturbation Algorithm (LPA) [9]. In order to guarantee differential privacy, the LPA generates the noise according to a Laplace distribution. $Lap(\lambda)$ denotes a random variable drawn from a Laplace distribution with probability density function (PDF): $\Pr[Lap(\lambda) = h] = \frac{1}{2\lambda} e^{-|h|/\lambda}$, where $Lap(\lambda)$ has zero mean and variance $2\lambda^2$. We denote the noisy and differentially private values as $\widetilde{X}_i = X_i(D) + Lap(\lambda)$ for $i = 1, 2, \ldots, n$. Since we have a series of eye movement observations, the final noisy eye movement observations are generated as follows.

$$\widetilde{X}^n = X^n(D) + Lap^n(\lambda) \tag{4}$$

where $Lap^n(\lambda)$ is a vector of n independent $Lap(\lambda)$ random variables and $X^n(D)$ is the eye movement observations without noise. The LPA algorithm is ϵ-differentially private for $\lambda = \dfrac{\Delta_1(X^n)}{\epsilon}$ [9].

Next, we define the error function that we use to measure the differences between original X^n and noisy $\widetilde{X}^n$ observations. For this purpose, we use the metric NMSE defined as

$$\text{NMSE} = \frac{1}{n}\sum_{i=1}^{n}\frac{(X_i - \widetilde{X}_i)^2}{\overline{X}\overline{\widetilde{X}}} \tag{5}$$

where

$$\overline{X} = \frac{1}{n}\sum_{i=1}^{n} X_i\,, \qquad \overline{\widetilde{X}} = \frac{1}{n}\sum_{i=1}^{n}\widetilde{X}_i. \tag{6}$$

We define the utility metric as

$$\text{Utility} = \frac{1}{\text{NMSE}}. \tag{7}$$

Since differential privacy is achieved by adding random noise to the data, there is a utility-privacy trade-off. If too much noise is introduced, the adversary will not be able to infer any information from the differentially private data. However, this might mean that the data are perturbed too much and no further analyses could be done due to reduced utility. When the eye tracking and movements data are considered, it is important to have high utility especially for further analyses.

4 METHODS

Standard mechanisms used for differential privacy are vulnerable to temporal correlations since the independent noise realizations that are added to temporally correlated data could be useful for adversaries. One straightforward approach is to decorrelate the data before adding the noise. However, a strong decorrelation can remove important eye movement patterns, which might be bad for further eye movement analyses by standard classification algorithms. On the other hand, many eye movement features are extracted by using time windows, as in [32], which makes the features highly correlated. Furthermore, another challenge is that the duration of eye tracking recordings could change depending on the skills or personalities of the users. The longer duration causes an increased query sensitivity, which means that higher amount of noise should be added to achieve differential privacy. In addition, when the data are correlated, as in [37], ϵ' is defined as the actual privacy metric that is obtained when considering the fact that correlation can be used to obtain more information about the differentially private data by filtering, instead of ϵ. In this work, we discuss and propose generic low-complexity methods to keep ϵ' low for eye movement feature signals. To deal with correlated eye movement feature data, we propose different methods: FPA, chunk-based FPA (C-FPA) for original signal and chunk-based FPA for difference based sequences (DC-FPA). The sensitivity of each eye movement feature signal is calculated by using the L_w-distance such that

$$\begin{aligned}\Delta_w^f(X^n) &= \max_{p,\,q}\left\| X^{n,(p,f)} - X^{n,(q,f)}\right\|_w \\ &= \max_{p,\,q}\sqrt[w]{\sum_{t=1}^{n}\left(\left|X_t^{(p,f)} - X_t^{(q,f)}\right|\right)^w}\end{aligned} \tag{8}$$

where $X^{n,(p,f)}$ and $X^{n,(q,f)}$ denote observation vectors for a feature f from two random participants p and q, n denotes the maximum length of the observation vectors, and $w \in \{1, 2\}$.

4.1 Fourier Perturbation Algorithm (FPA)

The main aim of the FPA is to represent a signal with a small number of transform coefficients such that the query sensitivity of the representative signal decreases. A smaller query sensitivity decreases the noise power required to make the noisy signal differentially private. In the FPA, the signal is transformed into the frequency domain by applying the standard Discrete Fourier Transform (DFT), which is commonly applied as a non-unitary transform. Furthermore, the frequency domain representation of a signal consists of less correlated transform coefficients as compared to the time domain signal due to the high decorrelation efficiency of the DFT. One therefore reduces the correlation between the eye movement feature signals by applying a Fourier transform. After applying the DFT, the noise sampled from the LPA is added to the first k elements of $DFT(X^n)$ that correspond to k lowest frequency components, denoted as $F^k = DFT^k(X^n)$. Once the noise is added, the remaining part (of size $n-k$) of the noisy signal $\widetilde{F}^k$ is zero padded. The obtained signal is denoted as $PAD^n(\widetilde{F}^k)$. Then, using the Inverse DFT (IDFT), the padded signal is transformed back into the time domain. We can show that ϵ-differential privacy is satisfied by the FPA for

$$\lambda = \frac{\sqrt{n}\sqrt{k}\Delta_2(X^n)}{\epsilon} \tag{9}$$

unlike the value claimed in [29], as observed independently in [19]. The procedure is summarized in Algorithm 1.

Algorithm 1: Fourier Perturbation Algorithm (FPA)

Inputs: X^n, λ, k
Output: $\widetilde{X}^n$
1) $F^k = DFT^k(X^n)$.
2) $\widetilde{F}^k = LPA(F^k, \lambda)$.
3) $\widetilde{X}^n = IDFT(PAD^n(\widetilde{F}^k))$.

Since not all coefficients are used, in addition to the perturbation error caused by the added noise, a reconstruction error caused by the lossy compression is introduced. It is important to determine the number of used coefficients k to minimize the total error. We discuss how we choose k values for FPA based methods in Section 4.4 in addition to shortcomings of this selection method.

4.2 Chunk-based FPA (C-FPA)

One of the drawbacks of directly applying the FPA to the eye movement feature signals is having large query sensitivities for each feature f due to long signal sizes. To solve this problem, [32] proposes to subsample the signal using non-overlapping windows, which simply means to remove many data points. While subsampling decreases the query sensitivity, it also decreases the amount of data. Instead of subsampling, we propose to split each signal into smaller chunks, and then we apply the FPA as discussed in Section 4.1 to each chunk. We choose the chunk sizes of, e.g., 64 and 128 since there exist divide-and-conquer type of tree based implementation algorithms for fast DFT calculations that can be applied when the transform size is a power of 2. When the signals are split into chunks, chunk level query sensitivities are calculated and used rather than the sensitivity of the whole sequence. Differential privacy for the complete signal is preserved by Theorem 2 since the chunks are non-overlapping. As the chunk size decreases, the chunk level sensitivity decreases as well as the computational complexity. However, the parameter ϵ' that is calculated according to the correlation level becomes larger with smaller chunk sizes due to the fact that correlations between

neighboring data elements are larger in an eye movement dataset. Therefore, it is important to obtain a good trade-off between computational complexity and correlations to determine the optimal chunk size.

4.3 Difference- and chunk-based FPA (DC-FPA)

To tackle temporal correlations, we propose to convert the original eye movement feature signals into difference signals where differences between consecutive eye movement feature observations are calculated as

$$\widehat{X}_t^{(f)} = \left\{X_t^{(f)} - X_{t-1}^{(f)}\right\}\Big|_{t=2}^{n} \quad \text{and} \quad \widehat{X}_1^{(f)} = X_1^{(f)}. \tag{10}$$

Using the difference signals denoted by $\widehat{X}^{n,(f)}$, we aim to decrease the correlations before applying a differential privacy method. We claim that the ratio $\frac{\epsilon'}{\epsilon}$ becomes lower in the difference based method as compared to the FPA method applied to the original dataset. To support this claim, we show that the correlations in the difference signals decrease significantly as compared to the original signal in Section 5.1. This results in lower ϵ' and better privacy for the same ϵ. The low level eye movement characteristics are preserved by the difference signals.

We propose to apply the difference based method together with the C-FPA. Therefore, the differences are calculated inside chunks. The first element of each chunk is preserved. Then, the FPA mechanism discussed in Section 4.1 is applied to the difference signals by using query sensitivities that are calculated based on differences and chunks. For each chunk, the noisy difference signals are propagated to obtain the final noisy signals. This mechanism is differentially private by Theorem 1 and overall procedure is summarized in Algorithm 2.

Algorithm 2: Difference- and chunk-based FPA (DC-FPA)

Inputs: X^n, λ, k
Output: $\widetilde{X}^n$
1) $\widehat{X}_t = \left\{X_t - X_{t-1}\right\}\Big|_{t=2}^{n}$ and $\widehat{X}_1 = X_1$.
2) $\widetilde{\widehat{X}}^n = FPA(\widehat{X}^n, \lambda, k)$.
3) $\widetilde{X_t} = \left\{\widetilde{\widehat{X}_t} + \widetilde{\widehat{X}_{t-1}}\right\}\Big|_{t=2}^{n}$ and $\widetilde{X}_1 = \widetilde{\widehat{X}_1}$.

Since Theorem 1 can be applied to the difference-based method when consecutive differences are assumed to be independent, which is a valid assumption for eye movement feature data as we illustrate below, there is also a trade-off between the chunk sizes and utility for the DC-FPA. If a large chunk size is chosen, then the total ϵ value could be very large, which reduces privacy. Therefore, we choose chunk sizes of, e.g., 32 and 64 for the DC-FPA, which are chosen because they provide better performance for the DC-FPA as compared to larger chunk sizes used for the C-FPA.

4.4 Choice of the number of transform coefficients used

All proposed methods require a selection of a value for k. A small k value increases the reconstruction error, while a very large k value results in an increase in the perturbation error. Therefore, it is important to find the best k value that minimizes the sum of the two errors. In this work, we compare a large set of possible k values to choose the best values.

We apply the differential privacy mechanisms that are discussed in Sections 4.1, 4.2, and 4.3 by using 100 noisy evaluations in order to find optimal k values applied to features or chunks. Optimal k values are the ones with the minimum NMSE for each chunk, eye movement feature, and document type. In a distributed setting, in which there are multiple parties that apply the same mechanisms, each user needs to know k values in advance. However, in a

centralized setting, where one server applies the differential privacy, it is crucial to choose the k values in a differentially private way as well. In order to evaluate the differential privacy in the eye tracking area while taking the temporal correlations into account, we focus on optimal k values for this work. One shortcoming of this approach is that the optimal k value compromises some information about the data, which leaks privacy [29]. Our observation is that the information leaked by optimizing the parameter k is negligible as compared to the privacy reduction due to high data correlation. Therefore, we illustrate only the results with optimal k values. For a differentially private selection of k, one can use the Sampling Perturbation Algorithm (SPA) proposed in [29].

5 EVALUATIONS

In this section, we describe the data and metrics we use to compare different methods. For each method, we average the results over 100 noisy evaluations. The differential privacy methods are applied separately to each document type (comic, newspaper, and textbook) on the dataset which is discussed in Section 5.1. Applying differential privacy separately to each document type preserves ϵ-differential privacy by Theorem 2. In addition, for the utility calculations and classifier training, the optimal k values are used.

5.1 Eye Tracking Dataset

We evaluate our methods by using the public eye tracking dataset from [32] dedicated to privacy-preserving eye tracking that is collected with Oculus DK2 VR headset and Pupil eye-tracking add-on [18] from 20 participants (10 female, 10 male) for a reading task of three different document types (a comic, newspaper, and textbook) in a VR environment. For each recording and participant, the dataset consists of 52 eye movement feature sequences that are computed with a sliding window size of 30 seconds and a step size of 0.5 seconds. The feature extraction process that is applied makes the extracted eye movement features independent of the eye tracker device used for data collection.

We first show how the data is correlated using correlation coefficients obtained from the eye movement feature data. Since there are 52 eye movement features, it is not feasible to show all of them. Therefore, we have chosen the feature called *ratio large saccade* for illustration. The correlation coefficients of *ratio large saccade* for three document types over a time difference Δt with respect to the signal samples at, e.g., the fifth time instance for eye movement features and differences for all participants are depicted in Figs. 1 (a) and (b), respectively.

Correlations between the difference signals are significantly smaller than the correlations between the original eye movement feature signals. Therefore, the DC-FPA is less vulnerable to privacy reduction due to temporal correlations, affecting the value of ϵ'. In addition, we observe that values of four features, namely minimum values of wordbook features from 1 to 4, are all zeros in the entire dataset. Therefore, we exclude them in the error and total ϵ calculations.

5.2 Utility Results

We now evaluate the utility given in (7). We apply our methods separately to different document types; therefore, we report the utility results separately as well. Furthermore, the utilities of FPA-based methods are calculated by using the optimal k values.

Remark. The contributions of each time instance to the query sensitivity are observed to be close for the eye movement dataset used. Therefore, by using k values such that $k \leq \sqrt{n}$, we assume that ϵ-differential privacy is preserved with $\lambda = \dfrac{\sqrt{k}\Delta_2(X^n)}{\epsilon}$ for this dataset. Furthermore, the actual privacy metric ϵ' rather than ϵ should be analyzed for evaluations, as discussed above.

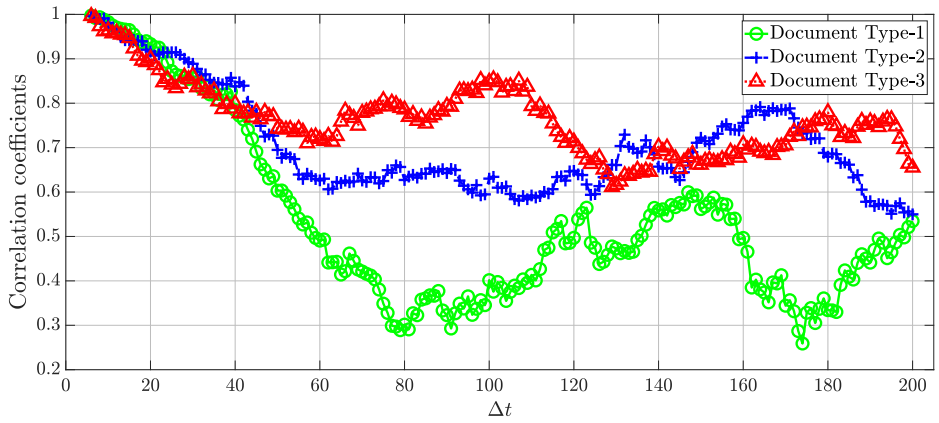

(a) Correlation coefficients of eye movement features.

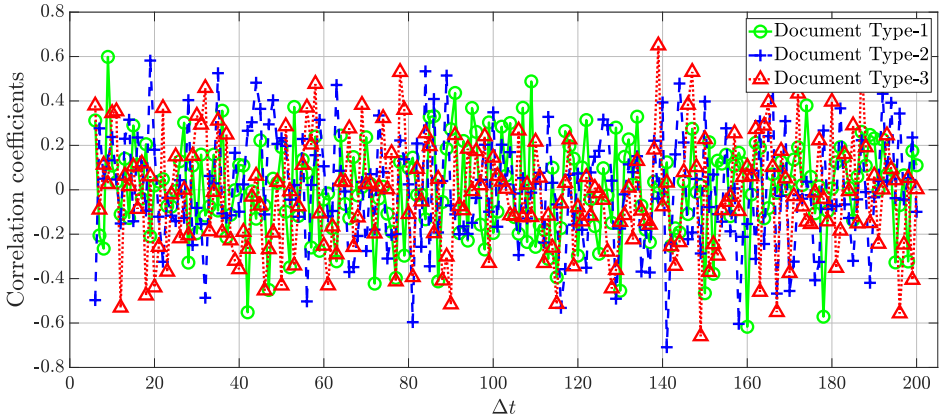

(b) Correlation coefficients of difference signals.

Fig. 1. Correlation coefficients of the eye movement feature *ratio large saccade* for three document types over a time difference of Δt (Each time instance step corresponds to 0.5s) with respect to the samples at the fifth time instance.

As we apply the proposed methods separately to each eye movement feature, we first calculate the mean utility of each eye movement feature and then calculate the average utility over all features. For the C-FPA and DC-FPA, we calculate the average over all eye movement features for each chunk. To obtain the global utility of the signal, we average the chunk level utilities. The utility results for an ϵ range of $[0.48, 2.4, 4.8, 24, 48]$ for LPA and FPA, C-FPA, and DC-FPA are given in Figs. 2 (a), (b), and (c), respectively.

While a high NMSE, i.e., low utility, does not necessarily mean that the model is completely wrong, higher utility means that the model would perform better than low utility in many tasks. Fig. 2 (a) shows that the LPA performs poorly for all document types and privacy levels. This is expected as the query sensitivity is large for the LPA as compared to FPA based methods. The C-FPAs depicted in Fig. 2 (b) for chunk sizes of 64 and 128 perform similarly. Since a higher

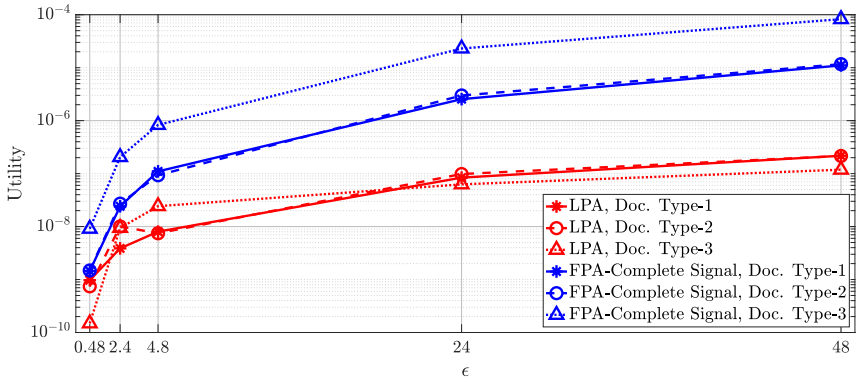

(a) Utility of the LPA and FPA.

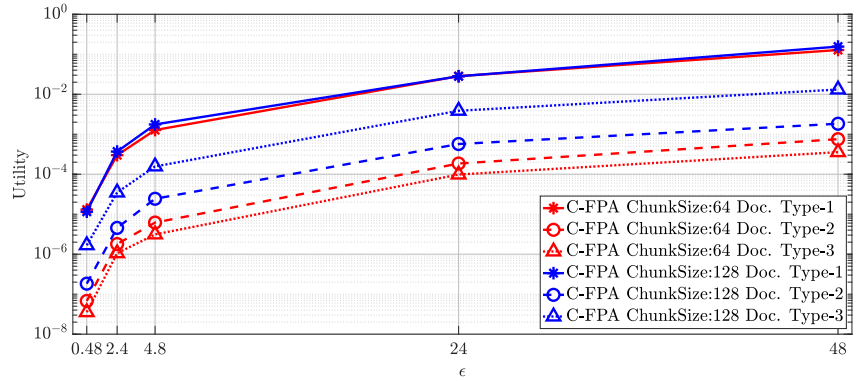

(b) Utility of the C-FPA.

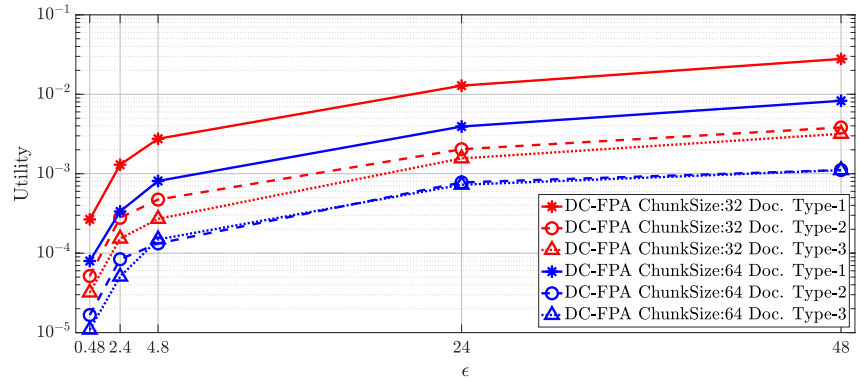

(c) Utility of the DC-FPA.

Fig. 2. Utility results for various ϵ values.

chunk size reduces the temporal correlations better, it is better to use a higher chunk size as long as the utilities are comparable. The C-FPAs with both chunk sizes outperform the LPA and FPA. In addition, applying FPA based methods to small chunks decreases the calculation complexity, which is another advantage of chunk-based methods.

The DC-FPA methods shown in Fig. 2 (c) outperform C-FPA methods in high privacy regions, i.e., for small ϵ. For less private regions, utility results are similar. Due to lower correlations in the DC-FPA, $\frac{\epsilon'}{\epsilon}$ is smaller for difference based methods as compared to the methods that use the original data.

5.3 Gender and Document Type Classifications

We evaluate gender and document type classification results for the differentially private data. For a fair comparison, we employ the same setup as in [32]. We train support vector machine (SVM) classifiers with radial basis function (RBF) kernel, bias parameter of $C = 1$, and automatic kernel scale in a leave-one-person-out cross-validation setup both with majority voting by summarizing classifications from different time instances for each participant and without majority voting. In the previous work, only majority voting based results are reported. However, a strong adversary might also train the models with different settings such as without majority voting. Therefore, we report the results for both cases. We normalize the training data to zero mean and unit variance, and apply the same parameters to the test data. Although we do not apply subsampling with a window size of 10 while generating the differentially private data, which is applied in [32], we use the subsampling in choosing training and testing data from each person and document type to have similar amount of data for training and testing to have a fair comparison. The results of the prediction accuracies of the methods for the gender and document type classification tasks using majority voting for various ϵ values are depicted in Figs. 3 (a) and (b), respectively.

While classification results could not be treated directly as the utility, they still provide insights into the usability of the differentially private data. In the majority voting setting, it is also possible to compare our classification results with the results in [32]. For the document type classification, all of our methods result in higher than the 0.33 guessing probability, even when $\epsilon < 1$. In addition, when ϵ is set to larger values such as 2.4 or 4.8, except DC-FPA methods all methods perform least at $\approx$ 55% and 65% accuracy, respectively. The method in the previous work only performs $\approx$ 40% for these privacy levels. Furthermore, when ϵ is set to larger values such as 24 or 48, our methods, except DC-FPA, result in accuracies greater than 85%, which significantly outperforms the previous work that has accuracies of $\approx 60 - 70\%$, for the same privacy regions. Although the DC-FPA methods do not outperform the previous work in terms of ϵ, they provide smaller ϵ' than the other methods. While we use the whole data rather than applying subsampling to generate differentially private data and handle the temporal correlations in a more robust way, we provide higher accuracies in the same privacy regions as compared to the results in [32].

Next, we analyze the gender classification results. While previously gender information is categorized as a feature that should be protected from adversaries, it could be interpreted in two ways. In order to compare our results with the previous work in [32] from the classification point of view, it could be argued that an adversary could not distinguish the genders when a certain amount of noise is added. In the majority voting setting, gender classification results are around 0.5, i.e., guessing probability, for $\epsilon \approx 4.8$ for the FPA and C-FPA methods. Since the classification results for the gender prediction are almost always below 0.5 for the DC-FPA methods, in the majority voting setting poor accuracy results are obtained. However, this effect is due to the majority voting setting and further validated by without majority voting setting as in Fig. 4 (a), since in this setting classification accuracies significantly approach to the guessing probability. In the previous work, the gender classification is performed around the guessing probability when $\epsilon \approx 15$. This shows

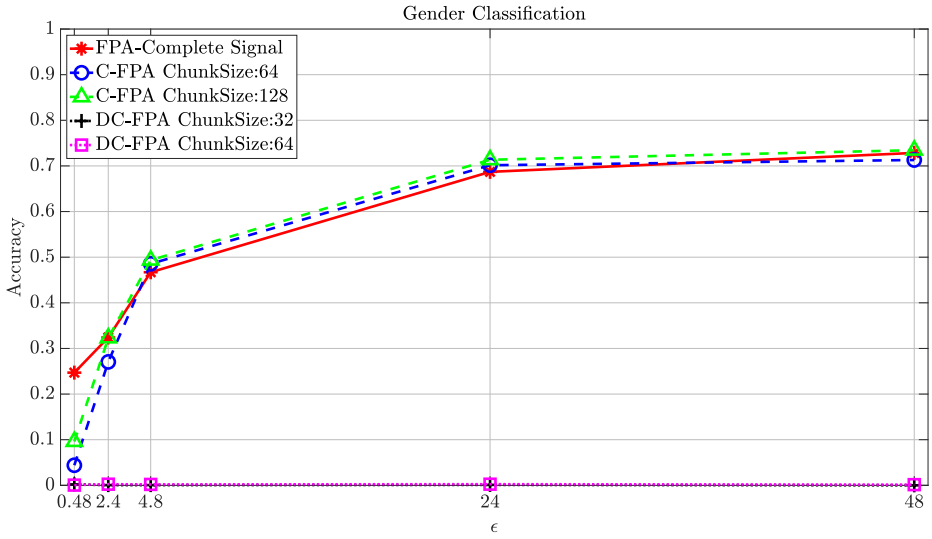

(a) Accuracy of gender classification with majority voting.

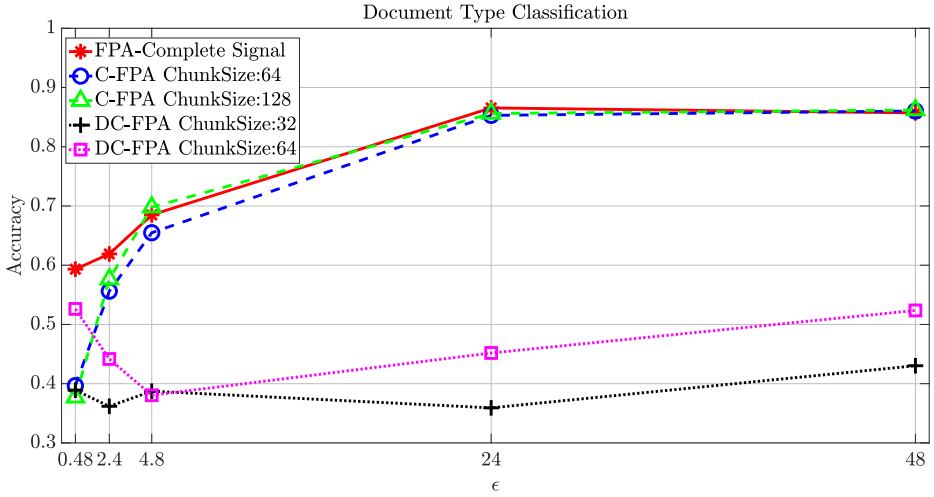

(b) Accuracy of document type classification with majority voting.

Fig. 3. Gender and document type classification accuracies with majority voting for various ϵ values.

that even if the gender classification is considered as a private task, we provide the privacy with lower ϵ values, which makes our methods more private. However, gender classification could be also interpreted as a part of utility. In this case, since we perform over the guessing probability for $\epsilon > 4.8$ with the C-FPA methods, it is reasonable to say that our methods are robust.

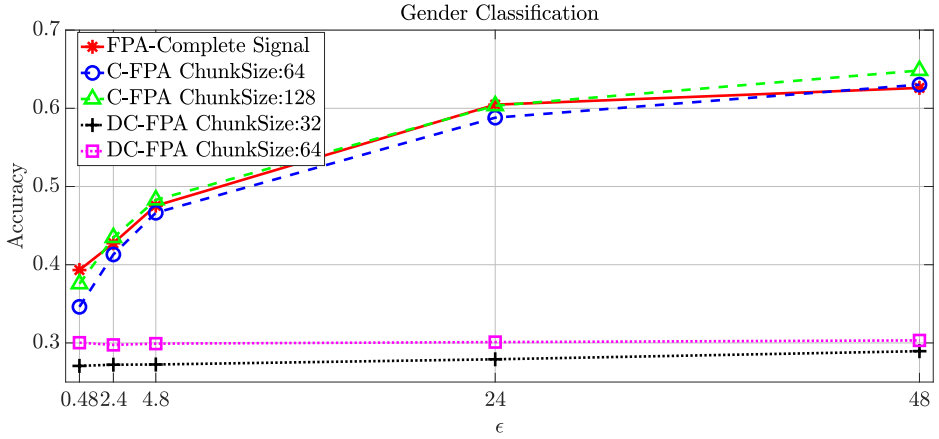

(a) Accuracy of gender classification without majority voting.

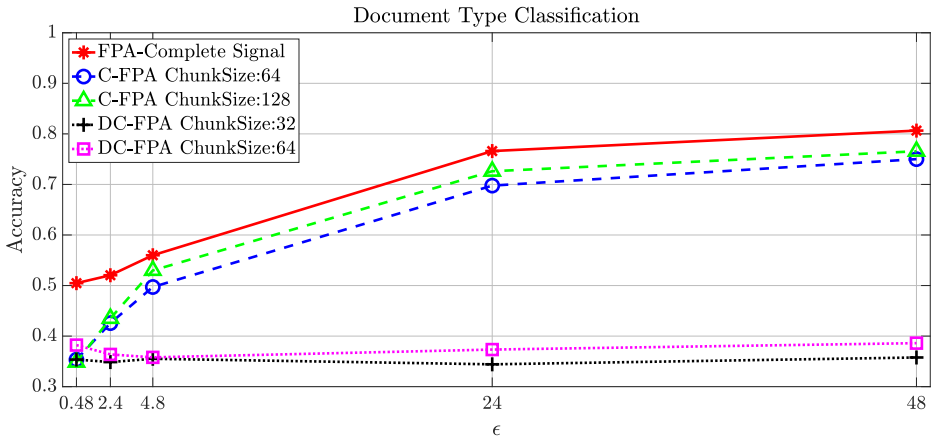

(b) Accuracy of document type classification without majority voting.

Fig. 4. Gender and document type classification accuracies without majority voting for various ϵ values.

We perform the same experiments without majority voting as well. Since an adversary does not have to follow one path to infer sensitive information, we also report the classification results without majority voting. While the similar trends are observed for some methods, it is visible that especially the gender classification results of the DC-FPA methods improve as compared to the majority voting setting. The prediction results for gender and document type classification without majority voting are shown in Figs. 4 (a) and (b), respectively.

For high privacy regions, i.e., for $\epsilon = 0.48$, the DC-FPA outperforms C-FPA in the document type classification task both for with and without majority voting setting as it is the case in the utility results in Section 5.2, and performs similarly to the FPA applied to the complete signal. It is clear that there are trade-offs which should be optimized to apply differential privacy in eye tracking area such as chunk sizes and correlations that affect ϵ'. We outperform previous work in the differentially private eye tracking literature also in terms of the classification accuracies.

6 CONCLUSION

We proposed different methods to achieve differential privacy by correcting, extending, and adapting the FPA method. Since eye movement features are correlated over time and are high dimensional, the standard privacy-preserving methods give low utility and are vulnerable to attacks. Taking these into consideration, we proposed privacy solutions for temporally correlated eye movement data. Our methods can be applied to the other human computer interaction data as well since the methods proposed are independent of the used data. We compared our methods with the previous work, and while taking care of the correlations robustly, we outperformed them in terms of utility and also the gender and document type classification accuracies. We also observed that to achieve higher privacy levels in the eye tracking and movements data, the high dimensional feature sequences that are obtained from the feature extraction processes should be transformed into more compact features since high dimensionality affects differential privacy negatively.

In future work, we will analyze the actual privacy metric ϵ' as the right privacy metric to compare different methods. Furthermore, since the eye movement features are extracted from fixed window and step sizes, these features are more correlated than standard signals. Therefore, we will use the actual measurements for ϵ' analyses with k values chosen in a private manner for the centralized setting.

ACKNOWLEDGMENTS

O. Günlü and R. F. Schaefer are supported by the German Federal Ministry of Education and Research (BMBF) within the national initiative for "Post Shannon Communication (NewCom)" under the Grant 16KIS1004. O. Günlü thanks Ravi Tandon for his useful suggestions. E. Bozkir thanks Martin Pawelczyk and Mete Akgün for useful discussions.

REFERENCES

[1] Tobias Appel, Christian Scharinger, Peter Gerjets, and Enkelejda Kasneci. 2018. Cross-subject Workload Classification Using Pupil-related Measures. In *ACM Symposium on Eye Tracking Research & Applications* (Warsaw, Poland). ACM, New York, NY, USA, 4:1–4:8. https://doi.org/10.1145/3204493.3204531

[2] Shlomo Berkovsky, Ronnie Taib, Irena Koprinska, Eileen Wang, Yucheng Zeng, Jingjie Li, and Sabina Kleitman. 2019. Detecting Personality Traits Using Eye-Tracking Data. In *ACM Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. ACM, New York, NY, USA, 221:1–221:12. https://doi.org/10.1145/3290605.3300451

[3] Efe Bozkir, David Geisler, and Enkelejda 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* (Barcelona, Spain) *(SAP '19)*. ACM, New York, NY, USA, Article 23, 5 pages. https://doi.org/10.1145/3343036.3343138

[4] 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. arXiv:1911.07936v2 [cs.CV] https://arxiv.org/abs/1911.07936v2

[5] Christian Braunagel, David Geisler, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2017. Online Recognition of Driver-Activity Based on Visual Scanpath Classification. *IEEE Intelligent Transportation Systems Magazine* 9, 4 (October 2017), 23–36. https://doi.org/10.1109/MITS.2017.2743171

[6] Andreas Bulling, Jamie A. Ward, Hans Gellersen, and Gerhard Troster. 2011. Eye Movement Analysis for Activity Recognition Using Electrooculography. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 4 (April 2011), 741–753. https://doi.org/10.1109/TPAMI.2010.86

[7] Nora Castner, Enkelejda Kasneci, Thomas Kübler, Katharina Scheiter, Juliane Richter, Thérése 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* (Warsaw, Poland) *(ETRA '18)*. ACM, New York, NY, USA, Article 39, 9 pages. https://doi.org/10.1145/3204493.3204550

[8] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. 2017. Collecting Telemetry Data Privately. In *International Conference on Neural Information Processing Systems* (Long Beach, California, USA). Curran Associates Inc., USA, 3574–3583.

[9] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.

[10] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3-4 (August 2014), 211–407. https://doi.org/10.1561/0400000042

[11] Simon Eberz, Kasper B. Rasmussen, Vincent Lenders, and Ivan Martinovic. 2016. Looks Like Eve: Exposing Insider Threats Using Eye Movement Biometrics. *ACM Trans. Priv. Secur.* 19, 1 (June 2016), 1:1–1:31. https://doi.org/10.1145/2904018

[12] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In *ACM SIGSAC Conference on Computer and Communications Security* (Scottsdale, Arizona, USA) *(CCS '14)*. ACM, New York, NY, USA, 1054–1067. https://doi.org/10.1145/2660267.2660348

[13] Onur Günlü. 2018. *Key Agreement with Physical Unclonable Functions and Biometric Identifiers*. Ph.D. Dissertation. TU Munich, Germany. published by Dr. Hut Verlag.

[14] Onur Günlü and Onurcan Iscan. 2014. DCT based ring oscillator Physical Unclonable Functions. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing* (Florence, Italy). IEEE, 8198–8201. https://doi.org/10.1109/ICASSP.2014.6855199

[15] Onur Günlü, Tasnad Kernetzky, Onurcan İşcan, Vladimir Sidorenko, Gerhard Kramer, and Rafael F. Schaefer. 2018. Secure and Reliable Key Agreement with Physical Unclonable Functions. *Entropy* 20, 5 (May 2018). https://doi.org/10.3390/e20050340

[16] Shoya Ishimaru, Kai Kunze, Koichi Kise, Jens Weppner, Andreas Dengel, Paul Lukowicz, and Andreas Bulling. 2014. In the Blink of an Eye: Combining Head Motion and Eye Blink Frequency for Activity Recognition with Google Glass. In *ACM Augmented Human International Conference* (Kobe, Japan). ACM, New York, NY, USA, 15:1–15:4. https://doi.org/10.1145/2582051.2582066

[17] Brendan John, Sanjeev Koppal, and Eakta Jain. 2019. EyeVEIL: Degrading Iris Authentication in Eye Tracking Headsets. In *ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado, USA). ACM, New York, NY, USA, Article 37, 5 pages. https://doi.org/10.1145/3314111.3319816

[18] Moritz Kassner, William Patera, and Andreas Bulling. 2014. Pupil: An Open Source Platform for Pervasive Eye Tracking and Mobile Gaze-Based Interaction. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication* (Seattle, Washington) *(UbiComp '14 Adjunct)*. ACM, New York, NY, USA, 1151–1160. https://doi.org/10.1145/2638728.2641695

[19] Georgios Kellaris and Stavros Papadopoulos. 2013. Practical differential privacy via grouping and smoothing, In VLDB Endowment. *Proc. VLDB Endow.* 6, 5, 301–312.

[20] Tomi Kinnunen, Filip Sedlak, and Roman Bednarik. 2010. Towards Task-independent Person Authentication Using Eye Movement Signals. In *ACM Symposium on Eye-Tracking Research & Applications* (Austin, Texas, USA). ACM, New York, NY, USA, 187–190. https://doi.org/10.1145/1743666.1743712

[21] Marion Koelle, Edgar Rose, and Susanne Boll. 2019. Ubiquitous Intelligent Cameras—Between Legal Nightmare and Social Empowerment. *IEEE MultiMedia* 26, 2 (April 2019), 76–86. https://doi.org/10.1109/MMUL.2019.2902922

[22] 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)* (Arlington, VA, USA). IEEE, 1–8. https://doi.org/10.1109/BTAS.2013.6712725

[23] Oleg V. Komogortsev, Sampath Jayarathna, Cecilia R. Aragon, and Mechehoul Mahmoud. 2010. Biometric Identification via an Oculomotor Plant Mathematical Model. In *ACM Symposium on Eye-Tracking Research & Applications* (Austin, Texas, USA). ACM, New York, NY, USA, 57–60. https://doi.org/10.1145/1743666.1743679

[24] Krzysztof Krejtz, Andrew T. Duchowski, Anna Niedzielska, Cezary Biele, and Izabela Krejtz. 2018. Eye tracking cognitive load using pupil diameter and microsaccades with fixed gaze. *PLOS ONE* 13, 9 (September 2018), 1–23. https://doi.org/10.1371/journal.pone.0203629

[25] Daniel J. Liebling and Sören Preibusch. 2014. Privacy Considerations for a Pervasive Eye Tracking World. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication* (Seattle, Washington, USA). ACM, New York, NY, USA, 1169–1177. https://doi.org/10.1145/2638728.2641688

[26] Ao Liu, Lirong Xia, Andrew Duchowski, Reynold Bailey, Kenneth Holmqvist, and Eakta Jain. 2019. Differential Privacy for Eye-tracking Data. In *ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado, USA) *(ETRA '19)*. ACM, New York, NY, USA, Article 28, 10 pages. https://doi.org/10.1145/3314111.3319823

[27] Frank D. McSherry. 2009. Privacy Integrated Queries: An Extensible Platform for Privacy-preserving Data Analysis. In *ACM SIGMOD International Conference on Management of Data* (Providence, Rhode Island, USA). ACM, New York, NY, USA, 19–30. https://doi.org/10.1145/1559845.1559850

[28] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust De-anonymization of Large Sparse Datasets. In *IEEE Symposium on Security and Privacy* (Oakland, CA, USA). IEEE, 111–125. https://doi.org/10.1109/SP.2008.33

[29] Vibhor Rastogi and Suman Nath. 2010. Differentially Private Aggregation of Distributed Time-series with Transformation and Encryption. In *ACM SIGMOD International Conference on Management of Data* (Indianapolis, Indiana, USA) *(SIGMOD '10)*. ACM, New York, NY, USA, 735–746. https://doi.org/10.1145/1807167.1807247

[30] Yosef Razin and Karen Feigh. 2017. Learning to Predict Intent from Gaze During Robotic Hand-Eye Coordination. In *AAAI Conference on Artificial Intelligence* (San Francisco, California, USA) *(AAAI'17)*. AAAI Press, 4596–4602.

[31] Julian Steil and Andreas Bulling. 2015. Discovery of Everyday Human Activities from Long-term Visual Behaviour Using Topic Models. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Osaka, Japan). ACM, New York, NY, USA, 75–85. https://doi.org/10.1145/2750858.2807520

[32] Julian Steil, Inken Hagestedt, Michael Xuelin Huang, and Andreas Bulling. 2019. Privacy-aware Eye Tracking Using Differential Privacy. In *ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado, USA) *(ETRA '19)*. ACM, New York, NY, USA, 27:1–27:9. https://doi.org/10.1145/3314111.3319915

[33] Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2019. PrivacEye: Privacy-preserving Head-mounted Eye Tracking Using Egocentric Scene Image and Eye Movement Features. In *ACM Symposium on Eye Tracking Research & Applications* (Denver, Colorado, USA) *(ETRA '19)*. ACM, New York, NY, USA, Article 26, 10 pages. https://doi.org/10.1145/3314111.3319913

[34] Peter M. van Leeuwen, Stefan de Groot, Riender Happee, and Joost C. F. de Winter. 2017. Differences between racing and non-racing drivers: A simulator study using eye-tracking. *PLOS ONE* 12, 11 (November 2017), 1–19. https://doi.org/10.1371/journal.pone.0186871

[35] Yasunori Yamada and Masatomo Kobayashi. 2018. Detecting mental fatigue from eye-tracking data gathered while watching video: Evaluation in younger and older adults. *Artificial Intelligence in Medicine* 91 (September 2018), 39 – 48. https://doi.org/10.1016/j.artmed.2018.06.005

[36] Yongtuo Zhang, Wen Hu, Weitao Xu, Chun Tung Chou, and Jiankun Hu. 2018. Continuous Authentication Using Eye Movement Response of Implicit Visual Stimuli. *ACM Interact. Mob. Wearable Ubiquitous Technology* 1, 4 (January 2018), 177:1–177:22. https://doi.org/10.1145/3161410

[37] Jun Zhao, Junshan Zhang, and H. Vincent Poor. 2017. Dependent Differential Privacy for Correlated Data. In *2017 IEEE Globecom Workshops (GC Wkshps)* (Singapore, Singapore). IEEE, 1–7. https://doi.org/10.1109/GLOCOMW.2017.8269219