Differential Privacy for Eye Tracking with Temporal Correlations

Efe Bozkir*†1 Onur Günlü*‡2 Wolfgang Fuhl†1 Rafael F. Schaefer‡2 Enkelejda Kasneci†1
[1]**Human-Computer Interaction, University of Tübingen**
[2]**Information Theory and Applications Chair, TU Berlin**

Abstract

New generation head-mounted displays, such as VR and AR glasses, are coming into the market with already integrated eye tracking and are expected to enable novel ways of human-computer interaction in many applications. However, since eye movement properties contain biometric information, privacy concerns have to be handled properly. Privacy-preservation techniques such as differential privacy mechanisms have recently been applied to the eye movement data obtained from such displays. Standard differential privacy mechanisms; however, are vulnerable to temporal correlations in the eye movement features. In this work, we propose a novel transform-coding based differential privacy mechanism to further adapt it to the statistics of eye movement feature data by comparing various low-complexity methods. We extent Fourier Perturbation Algorithm, which is a differential privacy mechanism, and correct a scaling mistake in its proof. Furthermore, we illustrate significant reductions in sample correlations in addition to query sensitivities, which provide the best utility-privacy trade-off in the eye tracking literature. Our results show significantly high privacy without loss in classification accuracies as well.

Introduction

Recent advances in the field of head-mounted displays (HMDs) and eye tracking enable easy access to pervasive eye trackers along with modern HMDs. Soon, the decrease in the cost of such devices might cause a mass consumption across different application domains such as gaming, entertainment, or education. Consequently, we expect a significant increase in the amount of eye movement data collected from the users. A large part of this data is indeed useful for personalized experience and user-adaptive interaction. In virtual and augmented reality (VR/AR) especially, it is possible to derive plenty of sensitive information about users from the eye movement data. For instance, it has been shown that eye tracking signals can be employed for activity recognition even in challenging everyday tasks (Steil and Bulling 2015; Braunagel et al. 2017; Ishimaru et al. 2014), to detect cognitive load (Appel et al. 2018; Krejtz et al. 2018), mental fatigue (Yamada and Kobayashi 2018), and many other user states. Similarly, assessment of situational attention (Bozkir, Geisler, and Kasneci 2019), expert-novice analysis in areas such as medicine (Castner et al. 2018) and sports (van Leeuwen et al. 2017), and personality traits and human intent during robotic hand-eye coordination can also be predicted based on eye movement features (Berkovsky et al. 2019; Razin and Feigh 2017). Additionally, eye movements are useful for early detection of anomias (Ungrady et al. 2019) and diseases (Fernández et al. 2015). Most importantly, eye movement data allow biometric authentication, which is considered to be a highly sensitive task (Günlü 2019). A task-independent authentication using eye movement features and Gaussian mixtures is, for example, thoroughly discussed in (Kinnunen, Sedlak, and Bednarik 2010). Additionally, biometric identification based on an oculomotor plant model and eye movements are introduced in (Komogortsev and Holland 2013; Komogortsev et al. 2010). In (Eberz et al. 2016) it is discussed that eye movement features can be used reliably also for authentication both in consumer level devices and various real world tasks, whereas (Zhang et al. 2018) shows that continuous authentication using eye movements is possible in VR headsets.

As biometric contents can be retrieved from eye movements, it is important to protect them against adversarial attacks. According to (Steil et al. 2019a), 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 considering that the usage of VR/AR devices with integrated eye trackers increases. As removing only the personal identifiers from data is not enough for anonymization due to linkage attacks (Narayanan and Shmatikov 2008), more sophisticated techniques for achieving user level privacy are necessary. Differential privacy (Dwork et al. 2006) is one effective solution, especially in the area of database applications. It protects user privacy by adding randomly generated noise for a given sensitivity and desired privacy parameter. However, high dimensionality of the data and temporal correlations can reduce utility and privacy, respectively. Since eye movement features are high dimensional, temporally correlated, and usually contain recordings with long durations, it is important to tackle utility and privacy problems simulta-

*Equally contributed
†efe.bozkir,wolfgang.fuhl,enkelejda.kasneci@uni-tuebingen.de
‡guenlue,rafael.schaefer@tu-berlin.de

neously. For eye movement data collected from HMDs or smart glasses, both local and global differential privacy can be applied. Local differential privacy adds user level noise to the data but assumes that the user sends data to a central data collector after adding local noise (Erlingsson, Pihur, and Korolova 2014; Ding, Kulkarni, and Yekhanin 2017). For this work, we consider global differential privacy, because there is a central user-level data collector and publisher in a VR/AR setting.

To apply differential privacy to the eye movement data, we evaluate the standard Laplacian Perturbation Algorithm (LPA) (Dwork et al. 2006) and Fourier Perturbation Algorithm (FPA) (Rastogi and Nath 2010). The latter is suitable for time series data such as the eye movement feature signals. 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 computational complexity. The difference-based method further decreases the temporal correlations between the eye movement features in addition to the decorrelation provided by the FPA that uses the discrete Fourier transform (DFT) as, e.g., in (Günlü and Iscan 2014; Günlü et al. 2018). The difference-based method provides a higher level of privacy since consecutive sample differences are observed to be less correlated than original consecutive data. Furthermore, we evaluate our methods using differentially private eye movement features in document type and gender classification, and privacy sensitivity classification tasks by using similar configurations to previous works in (Steil et al. 2019a,b). To generate differentially private eye movement data, we use the complete data instead of applying a subsampling step, used in (Steil et al. 2019a) to reduce the sensitivity and to improve the classification accuracies. In addition, the previous work in (Steil et al. 2019a) 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 (Dwork and Roth 2014). In eye movements, we are not interested in the "best" response but in the feature vector. Therefore, we apply the Laplacian mechanism. In summary, we are the first to propose differential privacy solutions for eye movements by taking the temporal correlations into account, which can help provide user privacy especially for HMD or smart glass usage.

Our main contributions are as follows. (1) We propose chunk-based and difference-based differential privacy methods for eye movement features to reduce query sensitivities, computational complexity, and temporal correlations. Furthermore, (2) we evaluate our methods on two publicly available eye movement datasets, i.e., MPIIDPEye (Steil et al. 2019a) and MPIIPrivacEye (Steil et al. 2019b), by comparing them with standard techniques such as LPA and FPA using the multiplicative inverse of the normalized mean square error (NMSE) as the utility metric. In addition, we evaluate document type and gender classification, and privacy sensitivity classification accuracies as classification metrics using differentially private eye movements in MPIIDPEye and MPIIPrivacEye datasets, respectively. Our results show significantly better performance as compared to previous works and are capable of handling correlated data and decreasing query sensitivities by dividing the data into smaller chunks.

Related Work

There are few works that focus on privacy-preserving eye tracking. (Liebling and Preibusch 2014) provides motivation as to why privacy considerations are needed for eye tracking data by focusing on gaze and pupillometry. Practical solutions are; therefore, introduced to protect user identity and sensitive stimuli based on a degraded iris authentication through optical defocus (John, Koppal, and Jain 2019) and an automated disabling mechanism for the eye tracker's ego perspective camera with the help of a mechanical shutter depending on the detection of privacy sensitive content (Steil et al. 2019b). Furthermore, a function-specific privacy model for privacy-preserving gaze estimation task and privacy-preserving eye videos by replacing the iris textures are proposed in (Bozkir et al. 2020) and (Chaudhary and Pelz 2020), respectively. For the user identity protection, works that focus on differential privacy are more relevant for us. Recently, standard differential privacy mechanisms are applied to eye movements in VR (Steil et al. 2019a) and heatmaps (Liu et al. 2019). These works do not address the effects of temporal correlations in eye movements over time in the privacy context. In the privacy literature, there are privacy definitions such as the Pufferfish mechanism (Kifer and Machanavajjhala 2014) or the Olympus framework (Raval, Machanavajjhala, and Pan 2019) for correlated data. These works, however, have different assumptions. For example, Pufferfish requires a domain expert to specify potential secrets and discriminative pairs, and Olympus models privacy and utility requirements as adversarial networks. As our focus is to protect user identity in the eye movements, we opt for differential privacy by discussing the effects of temporal correlations in eye movements over time and propose methods to reduce them.

Theoretical Background

Differential privacy uses a metric to measure the privacy risk for an individual participating in a database. Considering a dataset with weights of N people and a mean function, when an adversary queries the mean function for N people, the average weight over N people is obtained. After the first query, an additional query for $N-1$ people automatically leaks the weight of the remaining person. 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 randomly chosen individual participated in the dataset. The amount of noise added should be calibrated carefully since a high amount of noise might decrease the utility. We next define differential privacy.

Definition 1 ϵ-Differential Privacy (ϵ-DP) (Dwork et al. 2006). *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*

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

The variance of the added noise depends on the query sensitivity, which is defined as follows.

Definition 2 Query sensitivity (Dwork et al. 2006). *For a random query X^n and $w \in \{1,2\}$, the query 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 the L_w-distance is defined as

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

We list theorems that are used in the proposed methods.

Theorem 1 *Sequential Composition Theorem (McSherry 2009). 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 (McSherry 2009). 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.*

We define the Laplacian Perturbation Algorithm (LPA) (Dwork et al. 2006). 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 a 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 $\widetilde{X}^n = X^n(D) + Lap^n(\lambda)$, 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 $\lambda = \Delta_1(X^n)/\epsilon$ (Dwork et al. 2006).

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 normalized mean square error (NMSE) defined as

$$\text{NMSE} = \frac{1}{n}\sum_{i=1}^{n} \frac{(X_i - \widetilde{X}_i)^2}{\overline{X}\overline{\widetilde{X}}} \tag{4}$$

where

$$\overline{X} = \frac{1}{n}\sum_{i=1}^{n} X_i\,, \qquad \overline{\widetilde{X}} = \frac{1}{n}\sum_{i=1}^{n} \widetilde{X}_i. \tag{5}$$

We define the utility metric as

$$\text{Utility} = \frac{1}{\text{NMSE}}. \tag{6}$$

As differential privacy is achieved by adding random noise to the data, there is a utility-privacy trade-off. Too much noise will lead to high privacy; however, it might also result in poor analyses on the further tasks on eye movements. Therefore, it is important to find a good trade-off.

Methods

Standard differential privacy mechanisms are vulnerable to temporal correlations, since the independent noise realizations that are added to temporally correlated data could be useful for adversaries. However, decorrelating the data before adding the noise might remove important eye movement patterns and provide poor results in analyses. Many eye movement features are extracted by using time windows, as in (Steil et al. 2019a,b), which makes the features highly correlated. 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 amounts of noise should be added to achieve differential privacy. In addition, when the data is correlated, as in (Zhao, Zhang, and Poor 2017), ϵ' is defined as the actual privacy metric that is obtained considering the fact that correlations 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' small for eye movement feature signals. To deal with correlated eye movement feature data, we propose three different methods: FPA, chunk-based FPA (CFPA) for original signal, and chunk-based FPA for difference based sequences (DCFPA). The sensitivity of each eye movement feature signal is calculated by using the L_w-distance such that

$$\begin{aligned}\Delta_w^f(X^n) &= \max_{p,\,q} \Big|\Big| X^{n,(p,f)} - X^{n,(q,f)} \Big|\Big|_w \\ &= \max_{p,\,q} \sqrt[w]{\sum_{t=1}^{n} \Big(\Big| X_t^{(p,f)} - X_t^{(q,f)} \Big| \Big)^w}\end{aligned} \tag{7}$$

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

Fourier Perturbation Algorithm (FPA)

In the FPA, the signal is represented 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 Discrete Fourier Transform (DFT), which is commonly applied as a non-unitary transform. 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. Therefore, the correlation between the eye movement feature signals is reduced by applying the DFT. After 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 and denoted as $PAD^n(\widetilde{F}^k)$. Lastly, 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 $\lambda = \frac{\sqrt{n}\sqrt{k}\Delta_2(X^n)}{\epsilon}$ unlike the value claimed in (Rastogi and Nath 2010), as observed independently in (Kellaris and Papadopoulos 2013). The procedure is summarized in Algorithm 1. 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 below.

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) $\widetilde{X}^n = IDFT(PAD^n(\widetilde{F}^k))$.

Chunk-based FPA (CFPA)

One drawback of directly applying the FPA to the eye movement feature signals is large query sensitivities for each feature f due to long signal sizes. To solve this, (Steil et al. 2019a) proposes to subsample the signal using non-overlapping windows, which means removing many data points. While subsampling decreases the query sensitivities, it also decreases the amount of data. Instead, we propose to split each signal into smaller chunks and apply the FPA to each chunk so that complete data can be used. We choose the chunk sizes of 32, 64, and 128 since there are divide-and-conquer type tree-based implementation algorithms for fast DFT calculations 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 accounts for the sample correlations might increase with smaller chunk sizes because correlations between neighboring samples are larger in an eye movement dataset. Therefore, a good trade-off between computational complexity and correlations is needed to determine the optimal chunk size.

Difference- and chunk-based FPA (DCFPA)

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

$$\widehat{X}_t^{(f)} = \left\{ X_t^{(f)} - X_{t-1}^{(f)} \right\}\Big|_{t=2}^{n} \quad , \quad \widehat{X}_1^{(f)} = X_1^{(f)}. \tag{8}$$

Using the difference signals denoted by $\widehat{X}^{n,(f)}$, we aim to further decrease the correlations before applying a differential privacy method. We conjecture that the ratio ϵ'/ϵ decreases in the difference-based method as compared to the FPA method. To support this conjecture, we show that the correlations in the difference signals decrease significantly as compared to the original signals. This results in lower ϵ' and better privacy for the same ϵ. The difference based method is applied together with the CFPA. Therefore, the differences are calculated inside chunks. The first element of each chunk is preserved. Then, the FPA mechanism is applied to the difference signals by using query sensitivities calculated based on differences and chunks. For each chunk, noisy difference signals are aggregated to obtain the final noisy signals. This mechanism is differentially private by Theorem 1. Since Theorem 1 can be applied to the DCFPA 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 DCFPA. 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 32, 64, and 128 for the DCFPA as well for evaluation. The DCFPA is summarized in Algorithm 2.

Algorithm 2: DCFPA

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

Choice of the Number of Transform Coefficients

The proposed methods require a selection of a value for k. A small k value increases the reconstruction error, while a 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 aforementioned differential privacy mechanisms by using 100 noisy evaluations to find optimal k values applied to features or chunks. Optimal k values have the minimum NMSE for each chunk, eye movement feature, and document or recording type. In a distributed setting, each user needs to know k values in advance. However, in a centralized setting, it is crucial to choose the k values in a differentially private manner. 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 (Rastogi and Nath 2010). Our observation is that the information leaked by optimizing the parameter k is negligible as compared to the privacy reduction due

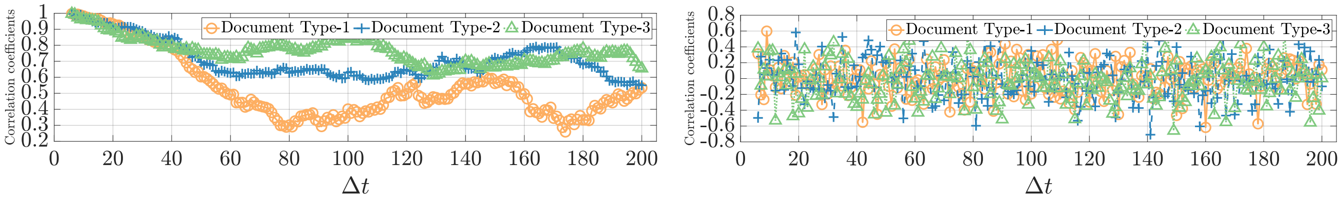

(a) Correlation coefficients of original signals in MPIIDPEye.

(b) Correlation coefficients of difference signals in MPIIDPEye.

Figure 1: Correlation coefficients of the feature *ratio large saccade* in MPIIDPEye dataset for three document types over a time difference of Δt (Each time step corresponds to 0.5s) w.r.t. the samples at the fifth time instance.

to correlated data. Thus, we illustrate the results with optimal k values.

Evaluations

This section discusses datasets, and evaluations using utility and classification metrics. The results are averaged over 100 noisy evaluations with the optimal k values in MATLAB.

Datasets

MPIIDPEye (Steil et al. 2019a): A publicly available eye movement dataset consisting of 60 recordings dedicated to privacy-preserving eye tracking that is collected in VR for a reading task of three document types (comic, newspaper, and textbook) from 20 (10 female, 10 male) participants. Each recording consists of 52 eye movement feature sequences computed with a sliding window size of 30 seconds and a step size of 0.5 seconds.

MPIIPrivacEye (Steil et al. 2019b): A publicly available eye movement dataset consisting of 51 recordings from 17 participants with 3 different sessions after each other with an head-mounted eye tracker and a field camera, which is similar to an AR setup. Each recording consists of 52 eye movement feature sequences computed with a sliding window size of 30 seconds and a step size of 1 second and each observation is annotated with binary privacy sensitivity levels of the scene that is being viewed. The dataset also consists of scene features extracted with CNNs. We do not evaluate the last part of the recording 1 of the participant 10, as the eye movement features are not available for this region. To detect the privacy level of the scene that is being viewed, we acknowledge that information about scene is very important (Orekondy, Schiele, and Fritz 2017); however, an individual's eye movements can improve the detection rate.

We first show the data correlation using correlation coefficients obtained from the eye movement features. Since there are 52 eye movement features in both datasets, it is not feasible to show them all. Thus, in the following we illustrate the correlation problem based on the feature called *ratio large saccade* in the MPIIDPEye dataset. The correlation coefficients of *ratio large saccade* for three document types over a time difference Δt w.r.t the signal samples at, e.g., the fifth time instance for original eye movement feature signals and difference signals for all participants are depicted in Figures 1 (a) and (b), respectively. As correlations between the difference signals are significantly smaller than correlations between the original eye movement feature signals, the DCFPA is less vulnerable to privacy reduction due to temporal correlations, thus affecting the value of ϵ'. Additionally, as all minimum values of wordbook features from 1 to 4 are zeros in both datasets, we exclude them from the utility and total ϵ calculations.

Utility Results

We evaluate the utility given in Equation (6) by applying our methods separately to different document and recording types; therefore, we report the utility results separately. As we apply the proposed methods separately to each eye movement feature, we first calculate the mean utility of each feature and then calculate the average utility over all features. The utility results for various ϵ values for aforementioned methods on MPIIDPEye and MPIIPrivacEye datasets are given in Figures 2 and 3, respectively.

While a high NMSE, i.e., low utility, does not necessarily mean that the model is completely useless, higher utility means that the model would perform more effectively than low utility in various tasks. The utility results of both evaluated datasets are similar. As the query sensitivities are lower in CFPA, utilities of CFPA are always higher than the utilities of the FPA as theoretically expected. DCFPA particularly outperforms other methods in the most private settings, namely in the lowest ϵ regions. When different chunk sizes are compared within the CFPA and DCFPA, different chunk sizes perform similarly for the CFPA method. For the DCFPA, there is a significant trend for better utilities when the chunk sizes are decreased. Since a higher chunk size reduces the temporal correlations better, it is ideal to use a higher chunk size if the utilities are comparable. While the LPA, namely the standard Laplacian mechanism of the differential privacy, is vulnerable to temporal correlations, our methods also outperform it in terms of utilities. In addition to high utilities, the calculation complexities are decreased with the CFPA and DCFPA which is another advantage of chunk-based methods.

Classification Accuracy Results

We evaluate document type and gender classification results for the MPIIDPEye and privacy sensitivity classification results for the MPIIPrivacEye by using differentially private data generated by the methods which handle temporal correlations in the privacy context. Instead of evaluating only Support Vector Machines (SVM) as in previous works (Steil

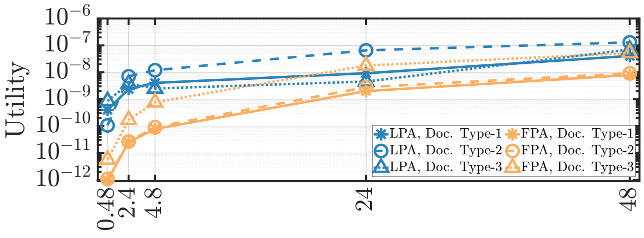

(a) Utility of the LPA and FPA for MPIIDPEye.

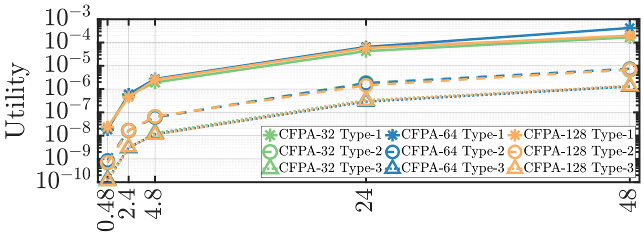

(b) Utility of the CFPA for MPIIDPEye.

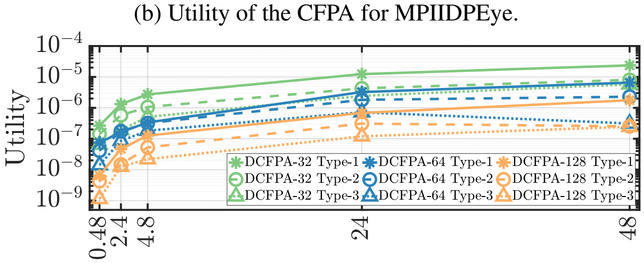

(c) Utility of the DCFPA for MPIIDPEye.

Figure 2: Utility results for MPIIDPEye dataset.

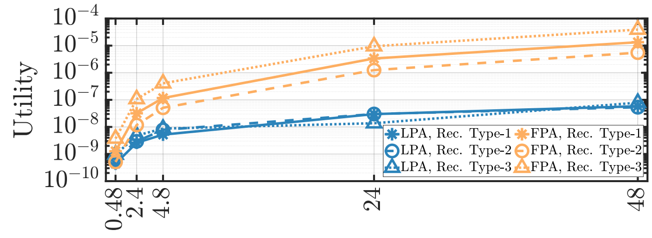

(a) Utility of the LPA and FPA for MPIIPrivacEye.

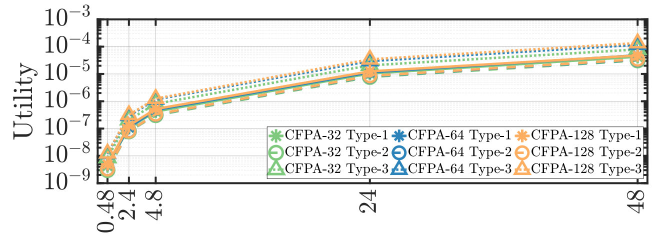

(b) Utility of the CFPA for MPIIPrivacEye.

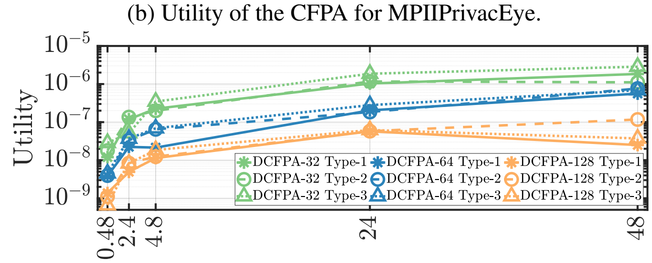

(c) Utility of the DCFPA for MPIIPrivacEye.

Figure 3: Utility results for MPIIPrivacEye dataset.

et al. 2019a,b), we evaluate a set of classifiers including SVMs, decision trees (DT), random forests (RF), and k-Nearest Neighbors (k-NN). We employ a similar setup as in (Steil et al. 2019a) with radial basis function (RBF) kernel, bias parameter of $C = 1$, and automatic kernel scale for the SVMs. For RFs and k-NNs, we use 10 trees and $k = 11$ with a random tie breaker among tied groups, respectively. 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 while generating the differentially private data, which is applied in (Steil et al. 2019a), we use subsampled data for training and testing with window sizes of 10 and 20 for MPIIDPEye and MPIIPrivacEye, respectively, to have a fair comparison and similar amount of data. All the classifiers are trained and tested in a leave-one-person-out cross-validation setup, which is considered as a more challenging but generic setup. For the MPIIDPEye, we evaluate results both with majority voting by summarizing classifications from different time instances for each participant and without majority voting. For the MPIIPrivacEye, it is not reasonable to use majority voting as each recording can include both privacy sensitive and non-sensitive stimuli.

While classification results cannot be treated directly as the utility, they provide insights into the usability of the differentially private data. We first evaluate document type classification task in the majority voting setting in Table 1 as it is possible to compare our results with the previous work (Steil et al. 2019a). As previous results quickly drop to the 0.33 guessing probability in high privacy regions, we significantly outperform them particularly with DCFPA and FPA with the accuracies over 0.60 and 0.85, respectively. In the less private regions towards $\epsilon = 48$, this trend still exists with the CFPA and FPA with accuracy results over 0.7 and 0.85. Chunk-based methods perform slightly worse than the FPA in the document type classifications even though the utility of the FPA is lower. We observe that the reading patterns are hidden easier with chunk-based methods; therefore, document type classification task becomes more challenging. This is especially validated with DCFPA methods using different chunk sizes, as DCFPA-128 outperforms smaller chunk-sized DCFPAs even though the sensitivities are higher. Therefore, we conclude that the differential privacy method should be selected for eye movements depending on the further task which will be applied.

Next, we analyze the gender classification results for MPIIDPEye. All methods are able to hide the gender information in the high privacy regions as it is already challenging to identify it with clean data as accuracies are ≈ 0.7 in

Document Type Classification Accuracies (k-NN|SVM|DT|RF)

Method	$\epsilon = 0.48$	$\epsilon = 2.4$	$\epsilon = 4.8$	$\epsilon = 24$	$\epsilon = 48$
FPA	$0.53 \vert 0.64 \vert 0.82 \vert \mathbf{0.88}$	$0.51 \vert 0.62 \vert 0.82 \vert \mathbf{0.87}$	$0.52 \vert 0.64 \vert 0.82 \vert \mathbf{0.87}$	$0.52 \vert 0.63 \vert 0.81 \vert \mathbf{0.87}$	$0.52 \vert 0.64 \vert 0.82 \vert \mathbf{0.88}$
CFPA-32	$0.39 \vert 0.37 \vert 0.45 \vert 0.43$	$0.40 \vert 0.39 \vert 0.45 \vert 0.44$	$0.40 \vert 0.44 \vert 0.46 \vert 0.44$	$0.59 \vert 0.58 \vert 0.55 \vert 0.60$	$\mathbf{0.71} \vert 0.69 \vert 0.66 \vert 0.66$
CFPA-64	$0.39 \vert 0.37 \vert 0.45 \vert 0.44$	$0.39 \vert 0.37 \vert 0.44 \vert 0.44$	$0.41 \vert 0.41 \vert 0.45 \vert 0.44$	$0.56 \vert 0.58 \vert 0.55 \vert 0.60$	$0.70 \vert 0.69 \vert 0.65 \vert 0.66$
CFPA-128	$0.37 \vert 0.33 \vert 0.45 \vert 0.45$	$0.37 \vert 0.32 \vert 0.44 \vert 0.45$	$0.38 \vert 0.36 \vert 0.44 \vert 0.45$	$0.52 \vert 0.55 \vert 0.51 \vert 0.56$	$0.68 \vert 0.68 \vert 0.63 \vert 0.65$
DCFPA-32	$0.50 \vert 0.37 \vert 0.47 \vert 0.43$	$0.51 \vert 0.35 \vert 0.47 \vert 0.42$	$0.49 \vert 0.34 \vert 0.47 \vert 0.43$	$0.49 \vert 0.37 \vert 0.45 \vert 0.44$	$0.49 \vert 0.36 \vert 0.47 \vert 0.44$
DCFPA-64	$0.60 \vert 0.44 \vert 0.43 \vert 0.41$	$0.56 \vert 0.35 \vert 0.42 \vert 0.40$	$0.58 \vert 0.41 \vert 0.43 \vert 0.41$	$\mathbf{0.60} \vert 0.43 \vert 0.44 \vert 0.42$	$0.60 \vert 0.41 \vert 0.44 \vert 0.43$
DCFPA-128	$\mathbf{0.64} \vert 0.47 \vert 0.46 \vert 0.48$	$\mathbf{0.62} \vert 0.43 \vert 0.45 \vert 0.45$	$\mathbf{0.68} \vert 0.50 \vert 0.45 \vert 0.46$	$0.57 \vert 0.45 \vert 0.45 \vert 0.47$	$0.60 \vert 0.41 \vert 0.45 \vert 0.45$

Table 1: Document type classification accuracies in MPIIDPEye using differentially private eye movement features with majority voting.

Gender Classification Accuracies (k-NN|SVM|DT|RF)

Method	$\epsilon = 0.48$	$\epsilon = 2.4$	$\epsilon = 4.8$	$\epsilon = 24$	$\epsilon = 48$
FPA	$0.42 \vert 0.28 \vert 0.42 \vert 0.37$	$0.43 \vert 0.29 \vert 0.42 \vert 0.39$	$0.42 \vert 0.27 \vert 0.44 \vert 0.38$	$0.44 \vert 0.28 \vert 0.40 \vert 0.40$	$0.43 \vert 0.30 \vert 0.43 \vert 0.39$
CFPA-32	$0.05 \vert 0.01 \vert 0.26 \vert 0.25$	$0.04 \vert 0.01 \vert 0.27 \vert 0.25$	$0.04 \vert 0.02 \vert 0.26 \vert 0.25$	$0.36 \vert 0.30 \vert 0.51 \vert 0.45$	$0.63 \vert 0.49 \vert \mathbf{0.69} \vert 0.54$
CFPA-64	$0.08 \vert 0.05 \vert 0.27 \vert 0.26$	$0.08 \vert 0.04 \vert 0.27 \vert 0.26$	$0.09 \vert 0.06 \vert 0.30 \vert 0.27$	$0.36 \vert 0.32 \vert 0.52 \vert 0.45$	$0.62 \vert 0.51 \vert \mathbf{0.68} \vert 0.54$
CFPA-128	$0.17 \vert 0.14 \vert 0.32 \vert 0.31$	$0.16 \vert 0.10 \vert 0.30 \vert 0.30$	$0.16 \vert 0.11 \vert 0.32 \vert 0.32$	$0.37 \vert 0.28 \vert 0.50 \vert 0.46$	$0.60 \vert 0.46 \vert \mathbf{0.67} \vert 0.53$
DCFPA-32	$0.03 \vert \approx 0 \vert 0.22 \vert 0.33$	$0.04 \vert \approx 0 \vert 0.23 \vert 0.32$	$0.04 \vert \approx 0 \vert 0.24 \vert 0.31$	$0.04 \vert \approx 0 \vert 0.22 \vert 0.31$	$0.04 \vert \approx 0 \vert 0.24 \vert 0.32$
DCFPA-64	$0.03 \vert \approx 0 \vert 0.28 \vert 0.33$	$0.04 \vert \approx 0 \vert 0.29 \vert 0.34$	$0.04 \vert \approx 0 \vert 0.28 \vert 0.34$	$0.03 \vert \approx 0 \vert 0.29 \vert 0.33$	$0.04 \vert \approx 0 \vert 0.29 \vert 0.33$
DCFPA-128	$0.08 \vert 0.01 \vert 0.32 \vert 0.34$	$0.08 \vert \approx 0 \vert 0.32 \vert 0.33$	$0.07 \vert \approx 0 \vert 0.33 \vert 0.35$	$0.08 \vert \approx 0 \vert 0.34 \vert 0.34$	$0.08 \vert \approx 0 \vert 0.32 \vert 0.33$

Table 2: Gender classification accuracies in MPIIDPEye using differentially private eye movement features with majority voting.

Privacy Sensitivity Classification Accuracies (k-NN|SVM|DT|RF)

Method	$\epsilon = 0.48$	$\epsilon = 2.4$	$\epsilon = 4.8$	$\epsilon = 24$	$\epsilon = 48$
FPA	$0.49 \vert 0.58 \vert 0.51 \vert 0.55$	$0.49 \vert 0.58 \vert 0.51 \vert 0.55$	$0.49 \vert 0.58 \vert 0.51 \vert 0.55$	$0.50 \vert 0.58 \vert 0.51 \vert 0.55$	$0.50 \vert 0.59 \vert 0.51 \vert 0.55$
CFPA-32	$0.55 \vert 0.59 \vert 0.52 \vert 0.56$	$0.55 \vert 0.58 \vert 0.52 \vert 0.56$	$0.55 \vert 0.58 \vert 0.52 \vert 0.56$	$0.56 \vert 0.58 \vert 0.53 \vert 0.57$	$0.58 \vert \mathbf{0.60} \vert 0.54 \vert 0.58$
CFPA-64	$0.55 \vert 0.58 \vert 0.52 \vert 0.56$	$0.55 \vert 0.58 \vert 0.52 \vert 0.56$	$0.55 \vert 0.58 \vert 0.52 \vert 0.56$	$0.56 \vert 0.58 \vert 0.53 \vert 0.57$	$0.58 \vert 0.59 \vert 0.54 \vert 0.58$
CFPA-128	$0.55 \vert 0.57 \vert 0.52 \vert 0.56$	$0.55 \vert 0.57 \vert 0.52 \vert 0.56$	$0.55 \vert 0.57 \vert 0.52 \vert 0.56$	$0.56 \vert 0.58 \vert 0.53 \vert 0.57$	$0.58 \vert 0.59 \vert 0.54 \vert 0.59$
DCFPA-32	$0.54 \vert \mathbf{0.59} \vert 0.52 \vert 0.56$	$0.55 \vert \mathbf{0.59} \vert 0.52 \vert 0.56$	$0.55 \vert \mathbf{0.59} \vert 0.52 \vert 0.56$	$0.54 \vert \mathbf{0.59} \vert 0.52 \vert 0.56$	$0.55 \vert 0.59 \vert 0.52 \vert 0.56$
DCFPA-64	$0.54 \vert 0.58 \vert 0.52 \vert 0.56$	$0.54 \vert 0.58 \vert 0.52 \vert 0.56$	$0.54 \vert 0.58 \vert 0.52 \vert 0.56$	$0.54 \vert 0.58 \vert 0.52 \vert 0.56$	$0.54 \vert 0.58 \vert 0.52 \vert 0.56$
DCFPA-128	$0.54 \vert 0.57 \vert 0.52 \vert 0.56$	$0.54 \vert 0.57 \vert 0.52 \vert 0.56$	$0.54 \vert 0.57 \vert 0.52 \vert 0.56$	$0.54 \vert 0.57 \vert 0.52 \vert 0.56$	$0.54 \vert 0.57 \vert 0.52 \vert 0.56$

Table 3: Privacy sensitivity classification accuracies in MPIIPrivacEye using differentially private eye movement features.

previous work (Steil et al. 2019a). While we obtain similar results compared to previous work for the gender classification task, the CFPA method is able to predict gender information correctly in the less private regions, namely $\epsilon = 48$, as it also has the highest utility values in these regions. The FPA applied to the complete signal and the DCFPA are not able to classify genders accurately. We observe that higher amount of noise that is needed by the FPA and removing the fine-grained "difference" information between eye movement observations with DCFPA are the reasons for hiding the gender information successfully in all privacy regions. Overall, the CFPA provides an optimal equilibrium between gender and document type classification success in the less private regions if gender information is not considered as a feature that should be protected from adversaries. Otherwise, all proposed methods are able to hide gender information from the data in the higher privacy regions as expected. Gender classification results are depicted in Table 2. Especially in some methods with k-NNs and SVMs, gender classification accuracies are close to zero because of the majority voting and it is validated by the results without majority voting in the Appendix.

For MPIIPrivacEye, we report privacy sensitivity classification accuracies using differentially private eye movements in the Table 3. The FPA performs worse than our methods. The DCFPA, particularly with the chunk size of 32, outperforms all other methods slightly in the higher privacy regions as it is also the case for the utility results. In the lower privacy regions, the CFPA performs the best with ≈ 0.60 accuracy. While having ≈ 0.60 accuracy in a binary classification problem does not form the best performance, according to the previous work (Steil et al. 2019b), privacy sensitivity classification using only eye movements with clean data in a person-independent setup only performs marginally higher than 0.60. Therefore, we show that even though we use differentially private data in the most private settings, we obtain similar results to the classification results using clean data. This means that differentially private eye movements can be used along with scene features for detecting privacy sensitive scenes in AR setups.

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, standard differential privacy methods provide low utility and are vulnerable to inference attacks. With this motivation, we proposed privacy solutions for temporally correlated eye movement data. Our methods can easily be applied to any other human-computer interaction data as well since they are independent of the used data. Our methods outperform state-of-the-art methods in terms of both utility and classification accuracies while taking care of the correlations robustly. In future work, we will analyze the actual privacy metric ϵ' with k values chosen in a private manner for the centralized differential privacy setting.

Ethics Statement

As head-mounted displays with integrated eye-tracking technology have found their way into many applications in daily life, it is possible to record a high amount of eye movement data. Apart from user assistive and comfort providing tasks, machines can identify biometric information using eye movement features. Differential privacy provides user privacy by adding randomly generated noise to the data. Especially with regard to data protection regulations, such as General Data Protection Regulation (GDPR) (EUd 2018), we foresee that manufacturers and users of not only head-mounted displays, but also any device that is integrated with eye trackers or sensors that collect temporally correlated personal data should benefit from this research. One disadvantage is that as a certain amount of noise is added to data for protection, for purposes such as gaze guidance or context sensitive aid, one may need more sophisticated approaches to deal with differentially private data. However, we also think that this would initiate new research directions in the field of human-computer interaction.

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

2018. The GDPR: New opportunities, new obligations.

Appel, T.; Scharinger, C.; Gerjets, P.; and Kasneci, E. 2018. Cross-subject Workload Classification Using Pupil-related Measures. In *ACM Symposium on Eye Tracking Research & Applications*, 4:1–4:8. New York, NY, USA. ISBN 978-1-4503-5706-7. doi:10.1145/3204493.3204531.

Berkovsky, S.; Taib, R.; Koprinska, I.; Wang, E.; Zeng, Y.; Li, J.; and Kleitman, S. 2019. Detecting Personality Traits Using Eye-Tracking Data. In *ACM Conference on Human Factors in Computing Systems*, CHI '19, 221:1–221:12. New York, NY, USA. ISBN 978-1-4503-5970-2. doi:10.1145/3290605.3300451.

Bozkir, E.; Geisler, D.; and Kasneci, E. 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, 23:1–23:5. New York, NY, USA. ISBN 978-1-4503-6890-2. doi:10.1145/3343036.3343138.

Bozkir, E.; Ünal, A. B.; Akgün, M.; Kasneci, E.; and Pfeifer, N. 2020. Privacy Preserving Gaze Estimation Using Synthetic Images via a Randomized Encoding Based Framework. In *ACM Symposium on Eye Tracking Research and Applications*, ETRA '20 Short Papers. New York, NY, USA: ACM. ISBN 9781450371346. doi:10.1145/3379156.3391364.

Braunagel, C.; Geisler, D.; Rosenstiel, W.; and Kasneci, E. 2017. Online Recognition of Driver-Activity Based on Visual Scanpath Classification. *IEEE Intelligent Transportation Systems Magazine* 9(4): 23–36. doi:10.1109/MITS.2017.2743171.

Castner, N.; Kasneci, E.; Kübler, T.; Scheiter, K.; Richter, J.; Eder, T.; Hüttig, F.; and Keutel, C. 2018. Scanpath Comparison in Medical Image Reading Skills of Dental Students: Distinguishing Stages of Expertise Development. In *ACM Symposium on Eye Tracking Research & Applications*, 39:1–39:9. New York, NY, USA. doi:10.1145/3204493.3204550.

Chaudhary, A. K.; and Pelz, J. B. 2020. Privacy-Preserving Eye Videos Using Rubber Sheet Model. In *ACM Symposium on Eye Tracking Research and Applications*, ETRA '20 Short Papers. New York, NY, USA: ACM. ISBN 9781450371346. doi:10.1145/3379156.3391375.

Ding, B.; Kulkarni, J.; and Yekhanin, S. 2017. Collecting Telemetry Data Privately. In *International Conference on Neural Information Processing Systems*, 3574–3583. USA: Curran Associates Inc.

Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In Halevi, S.; and Rabin, T., eds., *Theory of Cryptography*, 265–284. Berlin, Heidelberg: Springer Berlin Heidelberg.

Dwork, C.; and Roth, A. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9(3-4): 211–407. ISSN 1551-305X. doi:10.1561/0400000042.

Eberz, S.; Rasmussen, K. B.; Lenders, V.; and Martinovic, I. 2016. Looks Like Eve: Exposing Insider Threats Using Eye Movement Biometrics. *ACM Trans. Priv. Secur.* 19(1): 1:1–1:31. ISSN 2471-2566. doi:10.1145/2904018.

Erlingsson, U.; Pihur, V.; and Korolova, A. 2014. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In *ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, 1054–1067. New York, NY, USA. ISBN 978-1-4503-2957-6. doi:10.1145/2660267.2660348.

Fernández, G.; Manes, F.; Politi, L.; Orozco, D.; Schumacher, M.; Castro, L.; Agamennoni, O.; and Rotstein, N.

2015. Patients with Mild Alzheimer's Disease Fail When Using Their Working Memory: Evidence from the Eye Tracking Technique. *Journal of Alzheimer's disease: JAD* 50. doi:10.3233/JAD-150265.

Günlü, O. 2019. *Key Agreement with Physical Unclonable Functions and Biometric Identifiers*. Ph.D. thesis, TU Munich, Germany. Published by Dr. Hut Verlag.

Günlü, O.; and Iscan, O. 2014. DCT based ring oscillator Physical Unclonable Functions. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing*, 8198–8201. ISSN 2379-190X. doi:10.1109/ICASSP.2014.6855199.

Günlü, O.; Kernetzky, T.; İşcan, O.; Sidorenko, V.; Kramer, G.; and Schaefer, R. F. 2018. Secure and Reliable Key Agreement with Physical Unclonable Functions. *Entropy* 20(5). doi:10.3390/e20050340.

Ishimaru, S.; Kunze, K.; Kise, K.; Weppner, J.; Dengel, A.; Lukowicz, P.; and Bulling, A. 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*, 15:1–15:4. New York, NY, USA. ISBN 978-1-4503-2761-9. doi:10.1145/2582051.2582066.

John, B.; Koppal, S.; and Jain, E. 2019. EyeVEIL: Degrading Iris Authentication in Eye Tracking Headsets. In *ACM Symposium on Eye Tracking Research & Applications*, 37:1–37:5. New York, NY, USA. doi:10.1145/3314111.3319816.

Kellaris, G.; and Papadopoulos, S. 2013. Practical Differential Privacy via Grouping and Smoothing. *Proc. VLDB Endow.* 6(5): 301–312. ISSN 2150-8097. doi:10.14778/2535573.2488337.

Kifer, D.; and Machanavajjhala, A. 2014. Pufferfish: A Framework for Mathematical Privacy Definitions. *ACM Trans. Database Syst.* 39(1). ISSN 0362-5915. doi:10.1145/2514689.

Kinnunen, T.; Sedlak, F.; and Bednarik, R. 2010. Towards Task-independent Person Authentication Using Eye Movement Signals. In *ACM Symposium on Eye-Tracking Research & Applications*, 187–190. New York, NY, USA. ISBN 978-1-60558-994-7. doi:10.1145/1743666.1743712.

Komogortsev, O. V.; and Holland, C. D. 2013. Biometric authentication via complex oculomotor behavior. In *2013 IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, 1–8. doi:10.1109/BTAS.2013.6712725.

Komogortsev, O. V.; Jayarathna, S.; Aragon, C. R.; and Mahmoud, M. 2010. Biometric Identification via an Oculomotor Plant Mathematical Model. In *ACM Symposium on Eye-Tracking Research & Applications*, 57–60. New York, NY, USA. ISBN 978-1-60558-994-7. doi:10.1145/1743666.1743679.

Krejtz, K.; Duchowski, A. T.; Niedzielska, A.; Biele, C.; and Krejtz, I. 2018. Eye tracking cognitive load using pupil diameter and microsaccades with fixed gaze. *PLOS ONE* 13(9): 1–23. doi:10.1371/journal.pone.0203629.

Liebling, D. J.; and Preibusch, S. 2014. Privacy Considerations for a Pervasive Eye Tracking World. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, 1169–1177. New York, NY, USA. doi:10.1145/2638728.2641688.

Liu, A.; Xia, L.; Duchowski, A.; Bailey, R.; Holmqvist, K.; and Jain, E. 2019. Differential Privacy for Eye-tracking Data. In *ACM Symposium on Eye Tracking Research & Applications*, ETRA '19, 28:1–28:10. New York, NY, USA. ISBN 978-1-4503-6709-7. doi:10.1145/3314111.3319823.

McSherry, F. D. 2009. Privacy Integrated Queries: An Extensible Platform for Privacy-preserving Data Analysis. In *ACM SIGMOD International Conference on Management of Data*, 19–30. New York, NY, USA. doi:10.1145/1559845.1559850.

Narayanan, A.; and Shmatikov, V. 2008. Robust De-anonymization of Large Sparse Datasets. In *IEEE Symposium on Security and Privacy*, 111–125. doi:10.1109/SP.2008.33.

Orekondy, T.; Schiele, B.; and Fritz, M. 2017. Towards a Visual Privacy Advisor: Understanding and Predicting Privacy Risks in Images. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Rastogi, V.; and Nath, S. 2010. Differentially Private Aggregation of Distributed Time-series with Transformation and Encryption. In *ACM SIGMOD International Conference on Management of Data*, SIGMOD '10, 735–746. New York, NY, USA. ISBN 978-1-4503-0032-2. doi:10.1145/1807167.1807247.

Raval, N.; Machanavajjhala, A.; and Pan, J. 2019. Olympus: Sensor Privacy through Utility Aware Obfuscation. *Proceedings on Privacy Enhancing Technologies* 2019: 5–25. doi:10.2478/popets-2019-0002.

Razin, Y.; and Feigh, K. 2017. Learning to Predict Intent from Gaze During Robotic Hand-Eye Coordination. In *AAAI Conference on Artificial Intelligence*, 4596–4602.

Steil, J.; and Bulling, A. 2015. Discovery of Everyday Human Activities from Long-term Visual Behaviour Using Topic Models. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 75–85. New York, NY, USA. doi:10.1145/2750858.2807520.

Steil, J.; Hagestedt, I.; Huang, M. X.; and Bulling, A. 2019a. Privacy-aware Eye Tracking Using Differential Privacy. In *ACM Symposium on Eye Tracking Research & Applications*, ETRA '19, 27:1–27:9. New York, NY, USA. ISBN 978-1-4503-6709-7. doi:10.1145/3314111.3319915.

Steil, J.; Koelle, M.; Heuten, W.; Boll, S.; and Bulling, A. 2019b. PrivacEye: Privacy-preserving Head-mounted Eye Tracking Using Egocentric Scene Image and Eye Movement Features. In *ACM Symposium on Eye Tracking Research & Applications*, ETRA '19, 26:1–26:10. New York, NY, USA. ISBN 978-1-4503-6709-7. doi:10.1145/3314111.3319913.

Ungrady, M. B.; Flurie, M.; Zuckerman, B. M.; Mirman, D.; and Reilly, J. 2019. Naming and Knowing Revisited: Eyetracking Correlates of Anomia in Progressive Aphasia.

Frontiers in Human Neuroscience 13: 354. ISSN 1662-5161. doi:10.3389/fnhum.2019.00354.

van Leeuwen, P. M.; de Groot, S.; Happee, R.; and de Winter, J. C. F. 2017. Differences between racing and non-racing drivers: A simulator study using eye-tracking. *PLOS ONE* 12(11): 1–19. doi:10.1371/journal.pone.0186871.

Yamada, Y.; and Kobayashi, M. 2018. Detecting mental fatigue from eye-tracking data gathered while watching video: Evaluation in younger and older adults. *Artificial Intelligence in Medicine* 91: 39 – 48. ISSN 0933-3657. doi: https://doi.org/10.1016/j.artmed.2018.06.005.

Zhang, Y.; Hu, W.; Xu, W.; Chou, C. T.; and Hu, J. 2018. Continuous Authentication Using Eye Movement Response of Implicit Visual Stimuli. *ACM Interact. Mob. Wearable Ubiquitous Technology* 1(4): 177:1–177:22. ISSN 2474-9567. doi:10.1145/3161410.

Zhao, J.; Zhang, J.; and Poor, H. V. 2017. Dependent Differential Privacy for Correlated Data. In *2017 IEEE Globecom Workshops (GC Wkshps)*, 1–7. doi:10.1109/GLOCOMW.2017.8269219.

Supplementary Material

We report document type and gender classification results in MPIIDPEye dataset without majority voting in Tables 4 and 5, respectively.

	Document Type Classification Accuracies (k-NN|SVM|DT|RF)				
Method	$\epsilon = 0.48$	$\epsilon = 2.4$	$\epsilon = 4.8$	$\epsilon = 24$	$\epsilon = 48$
FPA	0.47|0.53|0.68|0.74	0.46|0.51|0.68|0.73	0.47|0.52|0.68|0.73	0.46|0.52|0.68|0.73	0.46|0.52|0.68|0.74
CFPA-32	0.34|0.35|0.36|0.38	0.34|0.35|0.36|0.38	0.34|0.36|0.36|0.38	0.39|0.45|0.38|0.42	0.47|0.53|0.44|0.49
CFPA-64	0.34|0.35|0.35|0.38	0.34|0.35|0.36|0.38	0.35|0.36|0.36|0.38	0.39|0.44|0.38|0.42	0.47|0.53|0.44|0.49
CFPA-128	0.34|0.34|0.36|0.39	0.34|0.34|0.36|0.39	0.34|0.34|0.36|0.39	0.38|0.42|0.37|0.42	0.46|0.51|0.43|0.49
DCFPA-32	0.36|0.35|0.36|0.37	0.36|0.34|0.36|0.37	0.35|0.34|0.36|0.37	0.36|0.35|0.36|0.37	0.35|0.34|0.36|0.38
DCFPA-64	0.38|0.37|0.35|0.37	0.37|0.35|0.35|0.37	0.37|0.36|0.35|0.37	0.37|0.36|0.35|0.37	0.37|0.36|0.35|0.37
DCFPA-128	0.40|0.38|0.36|0.38	0.39|0.37|0.35|0.38	0.41|0.39|0.36|0.38	0.38|0.37|0.35|0.38	0.39|0.37|0.35|0.37

Table 4: Document type classification accuracies in MPIIDPEye using differentially private eye movement features without majority voting.

	Gender Classification Accuracies (k-NN|SVM|DT|RF)				
Method	$\epsilon = 0.48$	$\epsilon = 2.4$	$\epsilon = 4.8$	$\epsilon = 24$	$\epsilon = 48$
FPA	0.47|0.41|0.48|0.44	0.47|0.42|0.47|0.45	0.47|0.41|0.48|0.45	0.47|0.41|0.47|0.45	0.48|0.42|0.48|0.45
CFPA-32	0.44|0.31|0.45|0.41	0.44|0.31|0.45|0.41	0.43|0.32|0.45|0.41	0.46|0.42|0.49|0.48	0.51|0.47|0.53|0.53
CFPA-64	0.44|0.35|0.45|0.41	0.44|0.35|0.45|0.41	0.44|0.35|0.46|0.42	0.46|0.43|0.49|0.43	0.51|0.48|0.54|0.53
CFPA-128	0.45|0.39|0.46|0.42	0.45|0.38|0.46|0.42	0.45|0.38|0.46|0.42	0.46|0.43|0.49|0.47	0.51|0.47|0.53|0.53
DCFPA-32	0.44|0.27|0.45|0.42	0.44|0.27|0.45|0.42	0.44|0.27|0.45|0.42	0.44|0.27|0.45|0.42	0.44|0.27|0.46|0.42
DCFPA-64	0.43|0.29|0.46|0.43	0.43|0.29|0.46|0.43	0.44|0.30|0.46|0.43	0.43|0.30|0.46|0.43	0.44|0.29|0.46|0.43
DCFPA-128	0.44|0.32|0.46|0.43	0.44|0.32|0.46|0.43	0.44|0.32|0.46|0.43	0.44|0.32|0.46|0.43	0.44|0.32|0.47|0.43

Table 5: Gender classification accuracies in MPIIDPEye using differentially private eye movement features without majority voting.