

For Your Eyes Only: Privacy-preserving eye-tracking datasets

Brendan David-John
brendanjohn@ufl.edu
University of Florida
USA

Kevin Butler
butler@ufl.edu
University of Florida
USA

Eakta Jain
ejain@cise.ufl.edu
University of Florida
USA

ABSTRACT

Eye-tracking is a critical source of information for understanding human behavior and developing future mixed-reality technology. Eye-tracking enables applications that classify user activity or predict user intent. However, eye-tracking datasets collected during common virtual reality tasks have also been shown to enable unique user identification, which creates a privacy risk. In this paper, we focus on the problem of user re-identification from eye-tracking features. We adapt standardized privacy definitions of k -anonymity and plausible deniability to protect datasets of eye-tracking features, and evaluate performance against re-identification by a standard biometric identification model on seven VR datasets. Our results demonstrate that re-identification goes down to chance levels for the privatized datasets, even as utility is preserved to levels higher than 72% accuracy in document type classification.

CCS CONCEPTS

• **Computing methodologies** → *Image processing*; • **Security and privacy** → *Privacy protections*; **Human and societal aspects of security and privacy**; • **Human-centered computing** → *Ubiquitous and mobile computing*; *Human computer interaction (HCI)*.

KEYWORDS

privacy, eye tracking, biometrics, re-identification, k -anonymity, plausible deniability

ACM Reference Format:

Brendan David-John, Kevin Butler, and Eakta Jain. 2022. For Your Eyes Only: Privacy-preserving eye-tracking datasets. In *2022 Symposium on Eye Tracking Research and Applications (ETRA '22)*, June 8–11, 2022, Seattle, WA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3517031.3529618>

1 INTRODUCTION

Re-identification attacks in literature have been extensively explored for social networks [Narayanan and Shmatikov 2009], location data [Primault et al. 2018], and medical data [El Emam et al. 2011]. Real-world re-identification attacks have been demonstrated to learn the medical prescriptions of a politician [Sweeney 2002] or reveal the Netflix preferences of half of a million users [Narayanan

and Shmatikov 2009]. As a result of the Netflix dataset attack, a woman sued the company over the risk that her leaked viewing patterns would reveal her sexual orientation to her family [Singel 2009]. There are an increasing number of algorithms that can authenticate a user based on eye movement data [George and Routray 2016; Lohr et al. 2021; Schröder et al. 2020; Sluganovic et al. 2018]. Numerous datasets of eye-tracking data for virtual reality (VR) applications are publicly available [David-John et al. 2021a; Emery et al. 2021; Hu et al. 2021; Sitzmann et al. 2018; Steil et al. 2019; Xu et al. 2018]. Taken together, this means that re-identification attacks using eye movements are not only plausible, but imminent.

Do people care? Surveys by Adams et al. [2018] and Steil et al. [2019] have established that both users and developers have privacy concerns over VR and eye-tracking data collection and how they are applied to make inferences about the user. For example, VR developers have cited that they are aware of privacy concerns for users and share their sentiments; however, most developers are not privacy experts and there is a lack of standards for how to address topics like ethics or privacy issues. For users, survey participants have indicated that they would be willing to accept beneficial VR applications that collect eye-tracking data if they are sharing the data with trusted governmental health agencies or with a university for research purposes. The same users also responded that they would not share their data publicly or with private services, unless there were constraints in place for how the data was being used.

Is regulation the answer? Privacy laws in certain regions are designed to protect traditional biometric identifiers, such as iris patterns and face scans [Heller 2020]. However, legal scholars have pointed out that privacy laws rarely hold up in court, and would not apply to behavioral data streams due to ambiguous wording over what is considered a biometric [Roberg-Perez 2016]. A lack of enforceable privacy laws and data release standards implies that VR platforms could store or sell identities through eye-tracking and behavioral data captured alongside demographics, which are typically used for personalized ads on the web [Datta et al. 2015].

Scope and contributions. In this paper we propose two novel adaptations of privacy mechanisms to achieve k -anonymity and plausible deniability (PD) guarantees for datasets of eye-tracking features. We compared our mechanisms against the previously established Exponential mechanism for DP. We found that our k -same-select sequence approach defended against re-identification and achieved superior utility in document type recognition ($\geq 72\%$).

2 RELATED WORK

Mechanisms that achieve formal privacy guarantees have been explored for protecting eye-tracking data against re-identification attacks for gaze samples [Li et al. 2020] and for features extracted

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ETRA '22, June 8–11, 2022, Seattle, WA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9252-5/22/06...\$15.00

<https://doi.org/10.1145/3517031.3529618>

Table 1: Privacy mechanisms for eye-tracking data with formal privacy guarantees. Shaded rows indicate our mechanisms.

| Mechanism | Guarantee | Data Type | Input to Mechanism | Adaption to Eye Tracking |
|---------------------------------------|------------------------|----------------------|--|--|
| Kaleido [Li et al. 2020] | ϵ, w, r -DP | Gaze Samples | Window of w gaze positions, current ROI radius r | Adapt spatial DP mechanism [Andrés et al. 2013] to incorporate a sequence [Kellaris et al. 2014] of gaze positions relative to dynamic ROIs in scene |
| Gaussian [Liu et al. 2019] | ϵ, δ -DP | Saliency Maps | User fixation map | Adapt DP noise mechanism [Dwork 2006] to protect fixation counts over image pixels |
| Exponential-DP [Steil et al. 2019] | ϵ -DP | Statistical Features | Gaze features extracted over window of time t | Adapt DP Noise mechanism [Dwork et al. 2014] applied to features independently |
| DCFPA [Bozkir et al. 2021] | ϵ -DP | Statistical Features | Gaze features extracted over window of time t | Adapt Fourier DP mechanism [Rastogi and Nath 2010] to include difference and chunking of sliding windows over time |
| <i>k</i> -same-select sequence (ours) | <i>k</i> -anonymity | Statistical Features | Gaze features extracted over window of time t | Randomly group features and apply <i>k</i> -same-select [Gross et al. 2005] over a sequence |
| Task-based Marginals (ours) | k, γ -PD | Statistical Features | Gaze features extracted over window of time t | Apply Marginals Generative Model and PD test [Bindschaedler et al. 2017] to each task |

from gaze data [Bozkir et al. 2021; Steil et al. 2019]. The un-shaded rows in Table 1 lists existing mechanisms that achieve formal privacy guarantees for eye-tracking data, type of input data, and how they were adapted to eye-tracking. The only formal privacy guarantee that has been explored is differential privacy (DP). While DP is popular in the privacy community due to the robust definition, there is an inevitable trade-off between increased DP privacy and lower data utility [Kifer and Machanavajjhala 2011].

We consider protecting eye-tracking datasets against re-identification attacks through alternative privacy guarantees. First, we explored *k*-anonymity to provide intuitive protection in that individual data cannot be distinguished from *k*-1 others. By adapting *k*-same-select [Gross et al. 2005], an upper bound on attack success is established while retaining utility. However, this approach releases *k* copies of the same data values. From an eye-tracking perspective, releasing duplicate data is not a satisfying solution. We shifted to considering PD, which extends a similar intuition for synthetic data. Synthetic data retains utility by reproducing characteristics of the original data. We explored guarantees specific to re-identification, and found superior utility with *k*-anonymity and that synthetic data has promise for preserving privacy in eye-tracking datasets.

The presented mechanisms are intended to be applied to datasets prior to their release. In contrast, for real-time systems, methods such as a privacy-preserving API and real-time perturbations [David-John et al. 2021a; Li et al. 2020] will enable platforms to share samples, features and gaze-based metrics with third-party applications.

3 METHODOLOGY

We conducted an evaluation of re-identification attacks on eye-tracking features and apply privacy mechanisms to protect identity. This section describes the protocol for re-identification attacks, privacy mechanisms for processing features, datasets included in the evaluation, and the approach used for biometric classification.

3.1 Threat Model

We assume that an adversary has access to a public eye-tracking dataset. The adversary trains an identification model to take eye-tracking feature vectors as input and output the associated identity.

Given new eye-tracking data, from playing a VR game for example, the adversary can use the trained model to guess at the identity of the player. The re-identification attack is successful if the correct identity is returned.

3.2 Proposed Solution

We propose two privacy mechanisms that can be applied to the eye tracking dataset prior to releasing it for public use. Thus, the adversary will train their model on privatized datasets. We assume that the adversary acquires un-privatized, i.e., raw data for the purposes of the re-identification attack, which is considered the test set. The privacy mechanisms are successful if they reduce the rate of re-identification to below chance levels.

3.3 Privacy Mechanisms

In this section, we contribute two privacy mechanisms, one that satisfies *k*-anonymity and one that satisfies plausible deniability. We provide pseudocode for ease of re-implementation and publicly release code for *k*-same-select sequence.¹ Both mechanisms are adaptations of prior work to consider eye-tracking features. For completeness, we provide pseudocode for our implementation of the DP-oriented mechanism defined by Steil et al. [2019].

k-same-select sequence. The *k*-same family of mechanisms [Gross et al. 2005; Newton et al. 2005] accomplish *k*-anonymity by first splitting individual data into groups of size *k*. Each group is averaged to produce a value which is then released *k* times in the released dataset. This enforces the upper bound on re-identification probabilities, as *k* of the identities from the original dataset will have equal contribution to the privatized data.

The implementation of *k*-same depends on the format of data being released. For example, *k*-same can be applied directly to face images by clustering and releasing averages [Newton et al. 2005]. For eye-tracking, the computed feature vectors are grouped and averaged to satisfy *k*-anonymity. We adapted the *k*-same-select mechanism by separately processing the sequence of feature vectors generated for each task in the dataset. Lines of code in black indicate

¹<https://doi.org/10.5281/zenodo.6463849>

the original k -same-select method and blue indicates our adapted steps. The data from all individuals are processed sequentially, i.e., the first feature vector of all individuals viewing a specific stimulus within a given task are randomly placed into groups of size k to compute average values for release. The mechanism assumes that there is data from at least k individuals available for grouping. The same groupings of individuals are used for each stimulus to achieve k -anonymity across the entire sequence of feature vectors. The adapted sequence mechanism is generalized by processing feature vectors in sequence; however, there is no guarantee that each individual had the same number of feature vectors per stimulus. Data are padded to repeat the last feature vector in the sequence for individuals with less features.

```

1: procedure  $k$ -SAME-SELECT SEQUENCE( $k$ , feature_data: structure indexing data by identity and task)
2:   for  $m = 1$  to  $num\_task$  do   ▶ Process features from each task
3:      $curr\_data \leftarrow feature\_data[m]$ 
4:      $G \leftarrow$  Randomize  $N$  individuals into  $H$  groups of size  $k$ 
5:     for  $i = 1$  to  $num\_feature\_vectors$  do   ▶ Loop over seq.
6:        $curr\_features \leftarrow curr\_data[i, :]$ 
7:        $avg\_features \leftarrow avg\_groups(curr\_features, G)$ 
8:        $curr\_data[i, :] \leftarrow avg\_features$ 
9:    $feature\_data[m] \leftarrow curr\_data$  ▶ Update task  $m$  features
return  $feature\_data$ 

```

Task-based Marginals Model (PD). Plausible deniability (PD) is not a condition of a privacy mechanism, but instead a privacy criterion that is checked before data is released [Bindschaedler et al. 2017]. Any number of approaches can be applied to generate data that satisfies PD. A generative model takes a raw feature vector as input and PD establishes that at least $k - 1$ other inputs from the original dataset could have plausibly generated the output synthetic feature. A parameter γ is used to control how close relative probabilities must be to be considered plausible, and k controls the number of features from the original dataset that have to pass the privacy test before synthetic data can be released. The formal definition and steps to implement the privacy test are detailed in the Supplementary Material.

To achieve PD we applied the Marginals approach with publicly available code [Bindschaedler et al. 2017].² Marginals builds a distribution of discrete values for each feature column and releases synthetic data by randomly sampling each feature independently. The learned feature distributions are representative of each task. Resulting distributions are used to synthesize data by task and retain utility. We adapted this approach by binning each continuous feature into $B = 30$ discrete buckets over the range of values (blue lines of code).

The generated synthetic feature vectors consist of discrete values corresponding to buckets that cover the range of feature values. We sample values between the minimum and maximum value range from the corresponding bucket from a random uniform distribution to map synthetic data back into continuous feature values. The synthetic dataset is stratified to contain the same number of feature vectors from each individual for each task as the original dataset. The PD guarantee differs from k -anonymity, in that PD guarantees $k - 1$ other **features** from the original dataset could have generated

the synthetic output, while k -anonymity guarantees that $k - 1$ other **individuals** could have generated a sequence of output features.

```

1: procedure TASK-BASED MARGINALS MODEL( $k$ ,  $\gamma$ ,  $B$ ,  $num\_samples$ , feature_data: structure indexing data by identity and task)
2:    $bin\_feature\_data \leftarrow BinData(feature\_data, B)$ 
3:   for  $m = 1$  to  $num\_task$  do   ▶ Process features from each task
4:      $M \leftarrow MarginalsDist(bin\_feature\_data[m])$ 
5:      $synth\_data \leftarrow Sample(M, num\_samples)$ 
6:      $private\_data \leftarrow PrivacyTest(synth\_data)$ 
7:      $bin\_feature\_data[m] \leftarrow private\_data$    ▶ Update task  $m$ 
8:    $feature\_data \leftarrow BinToContinuous(bin\_feature\_data)$ 
return  $feature\_data$ 

```

Exponential-DP Mechanism. The Exponential-DP noise mechanism was proven to be ϵ -DP by Steil et al. [2019] and applies to each individual feature in the feature set.³ Exponential noise is sampled independently for each feature vector and depends on the range of each feature and the task duration. The first step in applying Exponential-DP is to compute the range δ_i for each feature i as the maximum value minus the minimum value. The maximum number of feature vectors t_{max} from any individual during viewing is used for padding the data from other individuals. The last feature vector recorded for an individual is repeated to ensure that each individual has t_{max} total feature vectors. For each feature a value y is sampled from an Exponential distribution with a scale of $\frac{1}{\lambda}$, where $\lambda = \frac{\epsilon}{2 \cdot t_{max} \cdot \delta_i}$. The additive noise is then computed as $r = \pm \frac{\log_e(y)}{\lambda \cdot t_{max}}$ and the positive or negative sign is randomly assigned. Values of r are computed for every feature from the task, and are added to the original data to produce noisy feature vectors to release.

```

1: procedure EXPONENTIAL-DP( $\epsilon$ , feature_data: structure indexing data by identity and task)
2:    $\delta \leftarrow Range(feature\_data)$    ▶ Max value minus min
3:   for  $m = 1$  to  $num\_stimuli$  do ▶ Process features from each task
4:     Compute  $t_{max}$  for task  $m$  and pad individual data
5:      $\lambda \leftarrow \frac{\epsilon}{2 \cdot t_{max} \cdot \delta}$    ▶  $\lambda$  computed using  $\delta$ , and  $t_{max}$  from task
6:      $Exp \leftarrow Exponential(scales = \frac{1}{\lambda})$ 
7:     for  $i = 1$  to  $num\_feature\_vectors$  do   ▶ Loop over seq.
8:        $y \leftarrow Sample(Exp)$    ▶ Sample Exponential distribution
9:        $r \leftarrow \frac{\log_e(y)}{\lambda \cdot t_{max}}$    ▶ Compute additive noise value
10:     $feature\_data[m, i] \leftarrow feature\_data[m, i] \pm r$ 
return  $feature\_data$ 

```

3.4 Datasets

We evaluate the above detailed privacy mechanisms on publicly available VR datasets of eye-tracking features. The datasets vary based on the number of individuals, amount of data available, task being performed, and type of stimulus being viewed. Table 2 summarizes the characteristics of datasets included in our evaluation.

3.5 Feature Sets

Six of the datasets listed in Table 2 release raw gaze sample data, while MPIIDPEye included both raw samples and a set of pre-computed sliding windows of gaze-based features [Bulling et al. 2010]. To maintain consistency with past results from MPIIDPEye, we used their feature set in our analysis of this dataset. For all other

³Due to The Composition Theorem, Exponential-DP achieves a guarantee of ϵ times the number of features. For consistency with [Steil et al. 2019], we reference ϵ as the noise parameter for each feature, and not the composed guarantee.

²<https://vbinds.ch/node/69>

Table 2: Characteristics of VR eye-tracking datasets.

| Dataset | # Ppts. | Chance Rate | # Stim. | Data Per Ppt. | Stimuli Type | Task |
|------------------------------------|---------|--------------------|---------|---------------|--------------|---------------------------------------|
| MPIIDPEye [Steil et al. 2019] | 20 | 1/20 Ppts. = 5.0% | 3 | 30 mins | Documents | VR Reading |
| ET-DK2 [David-John et al. 2021a] | 18 | 1/18 Ppts. = 5.5% | 50 | 21 mins | 360° Images | Free Viewing |
| VR-Saliency [Sitzmann et al. 2018] | 130 | 1/130 Ppts. = 0.8% | 8 | 4 mins | 360° Images | Free Viewing |
| 360_em [Agtzidis et al. 2019b] | 13 | 1/13 Ppts. = 7.7% | 14 | 17 mins | 360° Videos | Free Viewing |
| VR-EyeTracking [Xu et al. 2018] | 43 | 1/43 Ppts. = 2.3% | 208 | Avg: 88 mins | 360° Videos | Free Viewing |
| OpenEDS [Emery et al. 2021] | 44 | 1/44 Ppts. = 2.3% | 2 | 10 mins | 3D Scene | Free Exploration |
| EHTask [Hu et al. 2021] | 30 | 1/30 Ppts. = 3.3% | 15 | 30 mins | 360° Videos | Free Viewing, Search, Saliency, Track |




Figure 1: Privacy evaluation for identification rate from eye-tracking features. Privatizing the dataset with our presented mechanisms lowers all identification rates to chance for $k = 8$ in k -same and Marginals, and $\epsilon = 2$ for Exponential-DP. Chance identification rates demonstrate that identity is protected within a group of individuals. The different datasets contain eye-tracking data on tasks performed within a variety of VR environments (reading documents, 360° images, 360° videos, and 3D rendered scenes). Chance rates (1/#Ppts.) vary for each dataset based on the number of identities, and are listed in Table 2.

datasets, we replicate the approach from David-John et al. [2021a] and extract features from fixation and saccade events detected using the I-S⁵T algorithm with default parameters [Agtzidis et al. 2019a]. The features extracted from fixation and saccades events leverage common statistics such as duration and amplitude, as well as the velocity and acceleration of gaze during the event [George and Routray 2016]. A feature set is generated for each type of event and a separate classification model is trained for each feature set.

3.6 Biometric Classifier

A Radial Basis Function (RBF) network is used to classify identity using feature vectors as input and is commonly used to identify users from eye-tracking data [David-John et al. 2021a; George and Routray 2016; Schröder et al. 2020]. An RBF network features a single hidden layer of nodes consisting of activation functions. The output of the activation functions is weighted to generate a probability that input is from each target class. The predicted class with the highest probability is considered the individual most likely to have produced the input feature vector, which is then used for biometric identification. Biometric identification relies on a set of features from an unknown individual viewing at least one stimulus. The feature vectors from all stimuli for an unknown individual are classified by the network, and the output scores are used to predict identity by averaging prediction scores.

As described in Section 3.5, the majority of datasets included in our evaluation use features extracted from both fixation and

saccade events, requiring an RBF network trained independently on both features [George and Routray 2016]. The output identification scores are first averaged within each type of event, then a final classification is made with a weighted average between fixation and saccade scores. A weight of 0.4 was applied to the fixation scores with a weight of 0.6 for saccade scores, as saccade features provided a slightly higher accuracy in user identification. For MPIIDPEye the prediction scores from all inputs within a task are simply averaged before classifying identity.

4 RESULTS

In this section, we present privacy and utility metrics to evaluate the implemented privacy mechanisms from Section 3.3 for each dataset listed in Table 2. We compared our proposed privacy mechanisms with Exponential-DP as an established approach for DP. Section 4.1 presents identification rates for each privacy mechanism using a biometric identification model trained on processed data and tested on the original data. Section 4.2 presents utility results for document type recognition on the MPIIDPEye dataset.

4.1 Biometric Identification

Re-identification risk for eye-tracking data is evaluated by splitting eye-tracking features into training sets processed by privacy mechanisms and testing sets of unmodified data. Identification rates higher than chance, which is one divided by the number of individuals in a dataset, indicate that there is risk of re-identification




Figure 2: Utility evaluation for accuracy of document type classification with an SVM model. Privatizing the dataset with our k -same mechanism retains the utility of the dataset for its intended application. In comparison, the Marginals Generative method does not retain utility above 53%, and the Exponential-DP mechanism rapidly leads to utility loss as we reach a parameter range of $\epsilon \leq 2$, where MPIIDPEye identification rates fell below chance.

from released data. Figure 1 presents the identification rates for each dataset and mechanism. The ET-DK2 dataset produced the highest identification rate of all datasets with 100% identification with the original data. All datasets produced identification rates higher than chance prior to privacy mechanisms being applied.

When privacy mechanisms were applied, the identification rates of all datasets dropped to chance. The Exponential-DP and Marginals approaches degraded the identification rates to chance across all parameter values. The only exception was MPIIDPEye for Exponential-DP, which required a parameter value of $\epsilon = 2$ for an identification rate of 6%, compared to a chance rate of 5%. k -same also reduces identification rates to chance, with a larger value of k needed to bring ET-DK2 to chance (5.6%). Our results suggest that privacy mechanisms protect against re-identification attacks on eye-tracking features using a standard biometric identification model.

4.2 Utility Evaluation

Releasing a privacy-preserving dataset that is useful relies on achieving a practical level of utility. We evaluated utility for each privacy mechanism applied to the MPIIDPEye dataset to classify document type being read using gaze features.

Steil et al. [2019] first evaluated MPIIDPEye using an SVM model to classify document type as either Comic, Newspaper, or Textbook. The SVM used an RBF kernel, bias parameter C set to one, and expressivity parameter γ set to one divided by the number of features. The model was trained on data from each individual during the first half of reading that was processed by the privacy mechanism, and tested on data from the second half. Figure 2 presents feature-level model accuracy results for each mechanism. Each plot demonstrates utility relative to the original data and chance rate of guessing (33%). We observed that the Exponential-DP mechanism reduced accuracy to chance, or near chance rates. For Exponential-DP, accuracy started at 80% for $\epsilon = 100$, and fell to chance at $\epsilon = 20$. For Marginals, a low level of utility was retained as accuracy remained near 53% for all parameters. The k -same approach was stable across parameter values, with slightly lower accuracy for higher levels of k . k -same across all parameters maintained performance greater than 72%. This level of accuracy would be practical

for an assistive reading interface that needs to identify the correct document type the majority of the time [Toyama et al. 2013].

5 CONCLUSION AND DISCUSSION

This paper addresses the open challenge of applying formal privacy definitions to behavioral data streams. Our work is the first to adapt the definitions of k -anonymity and PD to eye-tracking features. The definition of k -anonymity is intuitive as the theoretical risk of re-identification attacks are bounded above by $\frac{1}{k}$. The k -same-select sequence mechanism produced identification rates at chance while preserving model accuracy of 72% for document type classification. PD is a promising privacy criterion as it provides a clear interpretation with respect to re-identification, similar to k -anonymity; while using synthetic data to preserve privacy and retain utility. A Marginals mechanism for PD retains slight utility with an accuracy of 53% compared to a 33% guess rate. Deploying PD is computationally expensive, as a large-scale dataset of synthetic candidates are first generated before applying the privacy test. It took less than a minute to execute k -same and Exponential-DP, compared to roughly 30 minutes to generate and test synthetic data. Both k -same and Marginals mechanisms retain stable utility across their parameters, while the Exponential mechanism loses utility at the level of privacy needed for chance rates of identification.

Implications. The presented adaptations offer alternatives to DP, and demonstrate higher utility at chance rates for document type recognition. We recommend using k -same-select sequence for classification-based datasets to protect against re-identification as it is computationally efficient with an intuitive privacy guarantee.

Limitations. Our identification results were limited to an RBF network, although prior work explored random forest [Schröder et al. 2020], SVM [Miller et al. 2020], k -NNs [Bozkir et al. 2021] and deep network [Miller et al. 2021] models. In terms of DP, we only evaluated the Exponential-DP mechanism, although an alternative formulation of DP exists for time-series in the frequency domain [Bozkir et al. 2021]. While limitations impact the generalization of our empirical results, it does not impact the theoretical framing and comparison of privacy definitions.

Future Work. Our work provides motivation to adapt privacy guarantees to VR behavioral data in the form eye tracking. It would be useful to explore how well privacy methods preserve utility for other classification-based applications, such as intent prediction [David-John et al. 2021b]. Beyond exploring additional datasets and utilities, the field of eye-tracking privacy would benefit from further development of approaches related to PD. Such techniques can achieve an intuitive definition of privacy while preserving utility through synthetic data that appears real. Our proposed privacy mechanisms can also be applied to a breadth of mixed-reality sensors, including head and hand tracking, EEG, and EMG data.

ACKNOWLEDGMENTS

Eakta Jain acknowledges funding from NSF Award #2026540. Brendan David-John acknowledges funding from a Google PhD Fellowship and the National Science Foundation GRFP (Awards DGE-#1315138 and DGE-#1842473). Kevin Butler acknowledges funding from NSF Award CNS-#1815883 and CNS-#1562485, and AFOSR award FA#950-19-1-0169.

REFERENCES

- Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M Redmiles. 2018. Ethics emerging: the story of privacy and security perceptions in virtual reality. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. 427–442.
- Ioannis Agtzidis, Mikhail Startsev, and Michael Dorr. 2019a. 360-degree video gaze behaviour: A ground-truth data set and a classification algorithm for eye movements. In *Proceedings of the 27th ACM International Conference on Multimedia*. 1007–1015.
- Ioannis Agtzidis, Mikhail Startsev, and Michael Dorr. 2019b. A ground-truth data set and a classification algorithm for eye movements in 360-degree videos. *arXiv preprint arXiv:1903.06474* (2019).
- Miguel E Andrés, Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2013. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. 901–914.
- Vincent Bindschaedler, Reza Shokri, and Carl A Gunter. 2017. Plausible Deniability for Privacy-Preserving Data Synthesis. *Proceedings of the VLDB Endowment* 10, 5 (2017).
- Efe Bozkir, Onur Günlü, Wolfgang Fuhl, Rafael F Schaefer, and Enkelejda Kasneci. 2021. Differential privacy for eye tracking with temporal correlations. *Plos one* 16, 8 (2021), e0255979.
- Andreas Bulling, Jamie A Ward, Hans Gellersen, and Gerhard Tröster. 2010. Eye movement analysis for activity recognition using electrooculography. *IEEE transactions on pattern analysis and machine intelligence* 33, 4 (2010), 741–753.
- Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated Experiments on Ad Privacy Settings. *Proceedings on Privacy Enhancing Technologies* 2015, 1 (2015), 92–112.
- Brendan David-John, Diane Hosfelt, Kevin Butler, and Eakta Jain. 2021a. A privacy-preserving approach to streaming eye-tracking data. *IEEE Transactions on Visualization and Computer Graphics* 27, 5 (2021), 2555–2565.
- Brendan David-John, Candace Peacock, Ting Zhang, T Scott Murdison, Hrvoje Benko, and Tanya R Jonker. 2021b. Towards gaze-based prediction of the intent to interact in virtual reality. In *ACM Symposium on Eye Tracking Research and Applications*. 1–7.
- Cynthia Dwork. 2006. Differential Privacy. In *Automata, Languages and Programming*. Springer Berlin Heidelberg, 1–12. https://doi.org/10.1007/11787006_1
- Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (2014), 211–407.
- Khaled El Emam, Elizabeth Jonker, Luk Arbuckle, and Bradley Malin. 2011. A systematic review of re-identification attacks on health data. *PloS one* 6, 12 (2011), e28071.
- Kara J Emery, Marina Zannoli, James Warren, Lei Xiao, and Sachin S Talathi. 2021. OpenNEEDS: A Dataset of Gaze, Head, Hand, and Scene Signals During Exploration in Open-Ended VR Environments. In *ACM Symposium on Eye Tracking Research and Applications*. 1–7.
- Anjith George and Aurobinda Routray. 2016. A score level fusion method for eye movement biometrics. *Pattern Recognition Letters* 82 (2016), 207–215.
- Ralph Gross, Edoardo Airoldi, Bradley Malin, and Latanya Sweeney. 2005. Integrating utility into face de-identification. In *International Workshop on Privacy Enhancing Technologies*. Springer, 227–242.
- Brittan Heller. 2020. Watching Androids Dream of Electric Sheep: Immersive Technology, Biometric Psychography, and the Law. *Vanderbilt Journal of Entertainment & Technology Law* 23, 1 (2020), 1.
- Zhiming Hu, Andreas Bulling, Sheng Li, and Guoping Wang. 2021. EHTask: Recognizing User Tasks from Eye and Head Movements in Immersive Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* (2021).
- Georgios Kellaris, Stavros Papadopoulos, Xiaokui Xiao, and Dimitris Papadias. 2014. Differentially private event sequences over infinite streams. *Proceedings of the VLDB Endowment* 7, 12 (2014), 1155–1166.
- Daniel Kifer and Ashwin Machanavajhala. 2011. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. 193–204.
- Jingjie Li, Amrita Roy Chowdhury, Kassem Fawaz, and Younghyun Kim. 2020. Kaleido: Real-Time Privacy Control for Eye-Tracking Systems. In *29th USENIX Security Symposium (USENIX Security 20)*.
- Ao Liu, Lirong Xia, Andrew Duchowski, Reynold Bailey, Kenneth Holmqvist, and Eakta Jain. 2019. Differential privacy for eye-tracking data. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. ACM, 28.
- Dillon Lohr, Henry Griffith, and Oleg V Komogortsev. 2021. Eye Know You: Metric Learning for End-to-end Biometric Authentication Using Eye Movements from a Longitudinal Dataset. *arXiv preprint arXiv:2104.10489* (2021).
- Mark Roman Miller, Fernanda Herrera, Hanseul Jun, James A Landay, and Jeremy N Bailenson. 2020. Personal identifiability of user tracking data during observation of 360-degree VR video. *Scientific Reports* 10, 1 (2020), 1–10.
- Robert Miller, Natasha Kholgade Banerjee, and Sean Banerjee. 2021. Using Siamese Neural Networks to Perform Cross-System Behavioral Authentication in Virtual Reality. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, 140–149.
- Arvind Narayanan and Vitaly Shmatikov. 2009. De-anonymizing social networks. In *2009 30th IEEE symposium on security and privacy*. IEEE, 173–187.
- Elaine M Newton, Latanya Sweeney, and Bradley Malin. 2005. Preserving privacy by de-identifying face images. *IEEE transactions on Knowledge and Data Engineering* 17, 2 (2005), 232–243.
- Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, and Lionel Brunie. 2018. The long road to computational location privacy: A survey. *IEEE Communications Surveys & Tutorials* 21, 3 (2018), 2772–2793.
- Vibhor Rastogi and Suman Nath. 2010. Differentially private aggregation of distributed time-series with transformation and encryption. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*. 735–746.
- Sharon Roberg-Perez. 2016. The future is now: Biometric information and data privacy. *Antitrust* 31 (2016), 60.
- Christoph Schröder, Sahar Mahdie Klim Al Zaidawi, Martin HU Prinzer, Sebastian Maneth, and Gabriel Zachmann. 2020. Robustness of Eye Movement Biometrics Against Varying Stimuli and Varying Trajectory Length. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–7.
- Ryan Singel. 2009. Netflix spilled your Brokeback Mountain secret, lawsuit claims. *Threat Level (blog)*, *Wired* (2009).
- Vincent Sitzmann, Ana Serrano, Amy Pavel, Maneesh Agrawala, Diego Gutierrez, Belen Masia, and Gordon Wetzstein. 2018. Saliency in VR: How do people explore virtual environments? *IEEE Transactions on Visualization and Computer Graphics* 24, 4 (2018), 1633–1642.
- Ivo Služanović, Marc Roeschlin, Kasper B Rasmussen, and Ivan Martinović. 2018. Analysis of Reflexive Eye Movements for Fast Replay-Resistant Biometric Authentication. *ACM Transactions on Privacy and Security (TOPS)* 22, 1 (2018), 1–30.
- Julian Steil, Inken Hagedstedt, Michael Xuelin Huang, and Andreas Bulling. 2019. Privacy-Aware Eye Tracking Using Differential Privacy. In *11th ACM Symposium on Eye Tracking Research & Applications*. ACM.
- Latanya Sweeney. 2002. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, 05 (2002), 557–570.
- Takumi Toyama, Andreas Dengel, Wakana Suzuki, and Koichi Kise. 2013. Wearable reading assist system: Augmented reality document combining document retrieval and eye tracking. In *2013 12th International Conference on Document Analysis and Recognition*. IEEE, 30–34.
- Yanyu Xu, Yanbing Dong, Junru Wu, Zhengzhong Sun, Zhiru Shi, Jingyi Yu, and Shenghua Gao. 2018. Gaze Prediction in Dynamic 360° Immersive Videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5333–5342.