Privacy-preserving datasets of eye-tracking samples with applications in XR

Brendan David-John, *Member, IEEE*, Kevin Butler *Senior Member, IEEE* and Eakta Jain, *Member, IEEE*

Abstract—Virtual and mixed-reality (XR) technology has advanced significantly in the last few years and will enable the future of work, education, socialization, and entertainment. Eye-tracking data is required for supporting novel modes of interaction, animating virtual avatars, and implementing rendering or streaming optimizations. While eye tracking enables many beneficial applications in XR, it also introduces a risk to privacy by enabling re-identification of users. We applied privacy definitions of k-anonymity and plausible deniability (PD) to datasets of eye-tracking samples and evaluated them against the state-of-the-art differential privacy (DP) approach. Two VR datasets were processed to reduce identification rates while minimizing the impact on the performance of trained machine-learning models. Our results suggest that both PD and DP mechanisms produced practical privacy-utility trade-offs with respect to re-identification and activity classification accuracy, while k-anonymity performed best at retaining utility for gaze prediction.

Index Terms—Privacy, Eye Tracking, Re-identification, Biometrics

1 Introduction

Eye-tracking data presents a critical risk to privacy, as it captures sensitive information about the user based on where they look and introduces the risk of re-identification from captured data. Reducing the risk of re-identification from XR data was highlighted as the first recommendation of the IEEE Global Initiative on Ethics of Extended Reality report on XR and the Erosion of Anonymity and Privacy [47]: "XR stakeholders should actively develop and/or support efforts to standardize differential privacy and/or other privacy protocols that provide for the protection of individual identities and data." Eye trackers are among the XR sensors capable of accurate identification and recognition of users.

Eye-tracking data applied as a biometric is well studied both for iris recognition [32, 33] and gaze-based identification [20, 23, 24, 45, 59]. State-of-the-art eye movement biometrics can achieve an accuracy as high as 94% [59] and an Equal Error Rate of 2% [45], suggesting that, with high enough data quality, users are recognized as accurately as a four-digit pin with as little as five seconds of data [44].

Eye-tracking datasets are released publicly for research use or stored internally by XR companies to train proprietary models for product deployment. Datasets are anonymized by removing personal information such as names, locations, and dates of birth; however, they are still susceptible to re-identification attacks. A well-known example of a re-identification attack is from the Netflix Prize challenge [52]. Narayanan and Shmatikov took the released anonymous movie ratings combined with rental dates and matched them with public reviews from IMDB that were timestamped and linked to the user's real name. The risk of leaking users' identities and viewing patterns led to a lawsuit that claimed a woman's sexual orientation could be revealed to her family as a result of the attack [60]. In the above scenario, harm to individuals resulted from unauthorized disclosure of a sensitive attribute (i.e., viewing patterns that could reveal sexual orientation). Such risks are increasingly relevant for eye-tracking data, as gaze has the potential reveal age [67], sexual orientation [56], and personality traits [7].

- *Brendan David-John is an Assistant Professor at the Virginia Polytechnic Institute and State University. E-mail: bmdj@vt.edu.*
- *Dr. Kevin Butler is a Professor at the University of Florida. E-mail: butler@ufl.edu*
- *Dr. Eakta Jain is an Associate Professor at the University of Florida. E-mail: ejain@cise.ufl.edu*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

The current privacy recommendation for protecting eye-tracking samples against re-identification is through an API design that witholds gaze samples [16]. While privacy by limiting data access provides strong defense against re-identification, applications such as foveated rendering and gaze prediction can no longer be supported.

Current research in formal privacy guarantees for releasing gaze data have focused on differential privacy (DP) [11, 42, 43, 63]. Dwork proposed DP to release aggregate metrics while protecting individual data points with a strong privacy guarantee due to the formal bound on output distributions [18, 19]. DP is robust as it is able to protect privacy even in the worst-case scenario where an adversary has gained access to all other data points in the original dataset.

Research on privacy guarantees for novel data types and applications has introduced DP mechanisms into VR [51] and eye-tracking domains. DP mechanisms have been applied to different representations of gaze data, including saliency maps [43], eye images [34, 55], features extracted from eye movements [11, 63], and time series of gaze positions [42]. DP serves as the main benchmark for our proposed methods for gaze samples as DP is thus far the only formal guarantee for gaze samples. DP mechanisms have an inevitable impact on data utility [38], and identifying which applications are most affected or where alternative guarantees are a better fit can inform future privacy efforts for VR data. The goal of this work is to benchmark the privacy-utility trade-off of novel mechanisms that achieve k-anonymity and plausible deniability privacy guarantees for gaze samples applied to activity recognition and gaze prediction.

To our knowledge, this work is the first to provide a privacy guarantee alternative to DP for protecting eye-tracking sample datasets from re-identification attacks. The contributions of our work include:

- Novel privacy mechanisms for achieving k-anonymity and plausible deniability for generating synthetic eye-tracking sample datasets.
- Evaluation of privacy-utility trade-offs between k-anonymity, plausible deniability, and differential privacy when using eye-tracking datasets to train activity recognition and gaze prediction models.

In Section 2, we provide an overview of related work on privacy and eye-tracking data. Next, in Section 3, we motivate the need for privacy mechanisms by demonstrating re-identification attacks when eye-tracking data are paired with identifiers such as age and gender in public datasets. Then, Section 4 poses research questions, provides a threat model, defines the considered privacy guarantees, and discusses the implementation of the explored privacy mechanisms for sample data. Section 5 presents privacy and utility results for the mechanisms applied to two VR datasets (EHTask and DGaze). Last, Section 6 discusses limitations, takeaways, and directions for future work.

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

Mechanism	Guarantee	Data type	Input to mechanism	Adaption to eye tracking
Gaussian [43]	ε, δ-DP	Saliency maps	User fixation map	Adapt DP noise mechanism [17] to protect fixation counts over image pixels
Exponential-DP [63]	ε-DP	Statistical features	Gaze features extracted over window of time t	Adapt DP Noise mechanism [19] applied to features independently
DCFPA [11]	ε-DP	Statistical features	Gaze features extracted over window of time t	Adapt Fourier DP mechanism [54] to include difference and chunking of sliding windows
k-same-select sequence [15]	k-anonymity	Statistical features	Gaze features extracted over window of time t	Randomly group identities and apply k-same-select [26] over sequence of features
Task-based Marginals [15]	k, γ-PD	Statistical features	Gaze features extracted over window of time t	Apply Marginals Generative Model and PD test [10] to data from each task
Kalεido [42]	ε, w, r-DP	Gaze samples	Window of w gaze positions, spatial bound r	Adapt spatial DP mechanism [3] to incorporate a sequence [37] of gaze positions relative to ROIs detected in viewed content
k-same-synth (ours)	k-anonymity	Gaze samples	Gaze positions with event labels	Apply k-same-select sequence to parameters of models that generate event gaze positions
Event-synth-PD (ours)	k, γ-PD	Gaze samples	Gaze positions with event labels	Sample generative model for event gaze positions and apply PD Event Privacy test

2 Related Work

2.1 Privacy Guarantees for Eye-Tracking Data

Table 1 lists existing mechanisms that achieve formal privacy guarantees for eye-tracking data, the type of data input to the mechanism, and how the mechanism was adapted to eye tracking. The two most prominent data types listed in Table 1 are statistical features and gaze samples. Statistical features refer to statistics extracted during a fixed time window, such as the count of small, medium, or large amplitude saccades or the average fixation duration [12]. Alternatively, statistical features can be extracted from each individual fixation and saccade event, such as dwell time. Event-based features measure characteristics such as the maximum gaze velocity during a saccade or spatial dispersion during a fixation [24]. Statistical features summarize eye movement behavior, and are used for biometric identification of users [23] or classifying sensitive attributes [11, 63].

DP is a popular privacy guarantee that can protect released data from being used for identification or sensitive inferences. Specifically, DP bounds how much the output distribution of data changes in the case where any one feature vector in the dataset is omitted or included. DP methods for eye-tracking datasets add noise and release data in the same format as the input. The result of the DP guarantee means that the variation across individuals is reduced. Thus, identifying individual users or detecting sensitive traits becomes more difficult with privacy noise that masks individual differences. The privacy noise needed for DP also reduces the utility of the data by masking valuable insights.

Kalεido is the only existing mechanism for DP applied to sample data. DP in the context of spatial data ensures a probabilistic bound on how much output positions change within a spatial radius around the original data. The algorithm runs in real-time and allows for streaming samples with a DP guarantee; however, the context of the guarantee does not provide a theoretical bound on re-identification. It has been demonstrated that the amount of privacy noise added with kalεido reduces the risk of re-identification to chance for the 360_em dataset [42]. Yet, there is no analytical method to directly link the DP parameter ε with the theoretical risk of use re-identification.

Related work for eye-tracking feature datasets had introduced mechanisms that achieve k-anonymity and plausible deniability [15]. Compared to DP, both alternative guarantees retained higher utility for document type recognition when re-identification rates dropped to chance. The k-same-synth mechanism retained the highest classification accuracy while protecting privacy.

2.2 Alternative Privacy Guarantees

Many formal privacy guarantees exist to protect against different types of privacy risks. We pursued k-anonymity and k, γ-plausible deniability (PD) as alternatives to DP, as they directly protect against re-identification attacks. First, we explored k-anonymity to provide intuitive protection in that individual data cannot be distinguished from k-1 others. The k-same [26] approach is common to achieve k-anonymity for numerical data and works by averaging data together in groups of size k and releasing duplicate values. The duplicate values have equal contribution to the released data, establishing an upper bound of $\frac{1}{k}$ on the probability of individuals being re-identified. k-same is typically used to protect identity within facial images, as the numeric pixel values can easily be averaged across individuals. However, depending on the eye-tracking application, releasing duplicate data is not a satisfying solution.

Limiting output data to k duplicates to achieve k-anonymity led us to k, γ-PD, which extends a similar intuition applied to synthetically generated data [9, 10]. PD retains the intuition of the k parameter in terms of privacy for linking synthetic data to the real dataset. The γ parameter is used to threshold the probability that $k-1$ real data inputs could have generated the synthetic output before it can be released, allowing control over the level of privacy for data synthesized by a generative model. PD has been applied in the domain of spatial-temporal data in the form of location traces [9], motivating an application to spatial-temporal gaze data. Synthetic location traces retained utility for location-based services while protecting real individuals from being re-identified and leaking the specific location of their home, doctor's office, or work locations.

Adaptions of existing mechanisms for k-anonymity and k, γ-PD (Table 1) process feature data directly and allow for the protection of datasets that only release eye-tracking features extracted from raw data. The guarantees hold for the released feature data, as the only source of identification are the released feature values. In contrast, datasets of eye-tracking samples are difficult to protect against re-identification with a formal guarantee. The feature set an attacker may use for identification might not be known at the time of dataset release, preventing the privacy mechanism from providing a robust guarantee against future attacks. As described above, even the DP approach does not offer a direct theoretical guarantee against re-identification. Data release would require empirical analysis to determine which parameter values create data that is safe for release against a given feature set and model. To address limitations, we consider generative models that can synthesize gaze positions during the most common eye-movement events, fixation and saccades.

2.3 Synthesizing Gaze Data

Synthesizing eye-tracking data has been explored in the eye-tracking community to drive saliency-based applications [46, 68], and for training deep network models [40]. Generative models of gaze data are trained with the intention of deploying the model on new unseen inputs. For example, a deep model that predicts a fixation scanpath can take

as input an image and predict the most relevant regions to optimize during streaming. Deep synthesis models typically take the stimulus as input and predict the eye movement behavior of a viewer, which is considered synthetic data. In contrast, our proposed approach takes as input eye movements at the event level, synthesizing gaze positions during a saccade or fixation.

Past work modeling eye movements have developed simple [4,25,66] and complex models [36, 39] for events. Models vary in the number of parameters and whether they are based on heurisitics [6] or physical simulation [36]. Models based on statistical distributions are amenable to measuring the probabilities that relate individual samples with the distribution. Modeling the probability that an individual produced a set of samples for a particular event is key to preventing re-identification from extracted features.

We considered applications that process eye movement events or raw samples in our work. Such applications benefit from modeling eye-tracking data at a low level when compared to the high-level prediction models discussed above. For example, a real-time gaze model could predict where the user will look 100 milliseconds in the future. This task requires high utility of data within a short time scale, which is achieved by modeling each event detected in a sequence of eye-tracking data.

In the context of privacy, researchers have also turned to machine learning to learn how to transform real eye-tracking data with an autoencoder model to balance privacy and utility. Fuhl et al. [22] deployed a reinforcement learning model with privacy loss terms, optimizing the released data to achieve high privacy and utility for known classification models. The limitation of such an approach is the assumption that the attacker will use a similar classification model. However, with rapid advancements in deep learning biometrics, this assumption may not hold for the lifespan of the dataset [44]. Thus, formal privacy guarantees are the preferred approach when considering large-scale datasets and re-identification attacks.

3 Attack Scenario

A classic example of a re-identification attack is when the Governor of Massachusetts's medical prescriptions were leaked as a result of releasing gender, date of birth, and zip code [64]. Based on this example, consider a hypothetical dataset from a medical study that releases XR data publicly. The released data is de-identified by removing names and date of birth, but age and gender were retained along with eye-tracking data. The study was required to implement k-anonymity of k greater than or equal to four.[1] A k-anonymity guarantee can easily be achieved for the released age and gender data, however, if k-anonymity is not also achieved for the released eye-tracking data, then the risk of a successful re-identification attack is no longer bounded above by $\frac{1}{k}$. We demonstrate this risk with an example using the publicly available ET-DK2 [16] and 360_em [1] datasets that have recently been explored for re-identification [15].

In this scenario, it is assumed that the attacker can select the identities that match the demographics of their target and then train and apply a model to the subset of identities. For example, suppose only four identities in the dataset have an age between 18 and 20 and identify as Male. In that case, the attacker can train the model and predict which of the four identities is the target.

prototype of a gaze-based re-identification attack is conducted by combining the ET-DK2 and 360_em eye-tracking datasets with age and gender demographics. Age, gender, and a eye-tracking biometrics are all used for re-identification. The combined dataset in total includes 24 identities. A standard method of data generalization is used to achieve k-anonymity on age and gender labels by releasing ranges of values instead of exact values (see the Supplementary Material for the generalized k-anonymous groupings). Figure 1 demonstrates the success rate of re-identification attacks with and without the k-same-select sequence mechanism applied to the eye-tracking feature data [15].

[1] El Emam et al. [21] discussed k values for medical datasets, where k of three is a minimum, k of five is typical, and a k as large as fifteen is rare.

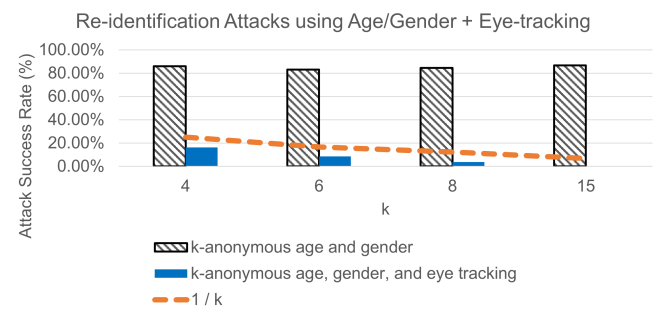

Fig. 1: Success rate of re-identification attacks using age, gender, and eye-tracking data. Bars with lines indicate results where age and gender are k-anonymous, while solid bars indicate results when eye-tracking data is also made k-anonymous. The orange dashed line plots $\frac{1}{k}$, the theoretical upper bound on re-identification.

A gaze-based biometric approach (see Sec. 4.5.1) was applied to perform the re-identification attack. Attack success remained above 80% for all values of k when only the age and gender demographics were k-anonymous. In contrast, attack success remained less than the theoretical $\frac{1}{k}$ bound when the eye-tracking data was also made to be k-anonymous prior to release.

In this example attack, a public dataset does not meet the k-anonymity privacy guarantee required by the research sponsor, impacting the researchers and institution that released the dataset. Furthermore, the scenario puts participants' privacy at risk, with successful re-identification attacks allowing the attacker to identify medical conditions or other sensitive information about victims.

4 Methods

The most pragmatic approach to prevent re-identification from gaze samples is through restricting access to raw data [16]. However, controlling access to raw data limits applications such as gaze prediction. We evaluate privacy mechanisms that reduce risk of user re-identification to explore the following questions,

RQ_1: Can synthetic data with formal privacy guarantees mitigate re-identification risk for datasets of gaze samples?

RQ_2: How do mechanisms that achieve k-anonymity and plausible deniability for gaze samples compare to a DP-based mechanism?

4.1 Threat Model

Assumptions for the considered re-identification attack include an attacker who has a target identity that they want to identify within the dataset. The attacker has read access to the public dataset. The attacker has access to eye-tracking data from the target performing the same task as the dataset. The attacker can then build a model trained on the public dataset that predicts which identity most closely matches the input data. If the prediction is correct, the target is successfully re-identified.

In this work, we considered a threat model where a privacy mechanism has processed the public dataset while the testing data used to re-identify individuals is unmodified. It is reasonable to assume that an attacker could gain access to raw tracking data through unauthorized code or by logging data streamed to third-party applications [65].

4.2 Privacy Definitions

This section defines three privacy definitions that can be applied to re-identification attacks. First, we discuss k-anonymity as the seminal definition of anonymity for a released dataset. Second, we present the definition of plausible deniability, which leverages the intuition of k-anonymity for synthetically generated data. Last, we provide the definition of ε-differential privacy.

4.2.1 k-anonymity

k-anonymity is a seminal definition of privacy within a dataset proposed by Samarati and Sweeney [58].

Definition 1 *k-anonymity*
Given a person-specific dataset D, a de-identified dataset D' is k-anonymized by privacy process $\mathscr{P} : D \mapsto D'$ if all released features $\Gamma_d = \mathscr{P}(\Gamma) \in D'$ cannot be recognized as Γ with probability $> \frac{1}{k}$.

A dataset has k-anonymity if the above condition is true for all unique combinations of feature values. There is no standard approach for determining k across fields, as the optimal value of k depends on the type of data and what probability of re-identification would make an individual feel safe from attacks.

4.2.2 Plausible Deniability

Plausible deniability (PD) was first defined by Bindschaedler and Shokri in the context of location traces [9] and later extended to general data formats [10]. PD prevents re-identification by utilizing the generation of synthetic data to achieve privacy. A synthetic dataset is released that captures the original characteristics without leaking the identity of those that contributed to the original dataset. PD provides a guarantee that there are at least k individual records that could have plausibly generated a synthetic data output.

PD has two privacy parameters: k, an integer greater than or equal to one, and γ, a real number greater than or equal to one.

Definition 2 *Plausible Deniability*
For any dataset D where $|D| \geq k$, and any record y generated by a probabilistic generative model $\mathbf{M}$ such that $y = \mathbf{M}(d_1)$ for $d_i \in D$, we state that y is releasable with (k,γ)-plausible deniability if there exist at least $k-1$ unique records $d_2, ..., d_k \in D \setminus \{d_1\}$, such that

$$\gamma^{-1} \leq \frac{Pr\{y = \mathbf{M}(d_i)\}}{Pr\{y = \mathbf{M}(d_j)\}} \leq \gamma$$

where $k \geq 1$ is an integer and $\gamma \geq 1$ is a real number.

Large values of k and values of γ that are closer to one imply higher privacy. Privacy-preserving datasets are generated by only releasing synthetic records y if they pass the PD privacy test:

1. Let i ≥ 0 be the only integer that fits the inequality $\gamma^{-i-1} < Pr\{y = \mathbf{M}(d)\} \leq \gamma^{-i}$
2. Let k' be the count of records $d_a \in D$ such that $\gamma^{-i-1} < Pr\{y = \mathbf{M}(d_a)\} \leq \gamma^{-i}$
3. If $k' \geq k$: return PASS, else return FAIL

Implementing the Privacy Test requires a method to compute probability values of the form $Pr\{y = \mathbf{M}(d_i)\}$ that represent the probability that the mechanism $\mathbf{M}$ would generate the synthetic output y for a given input. PD is intuitive against re-identification in terms of k, similar to k-anonymity. Using synthetic data achieves privacy while retaining data utility if the generated data captures the characteristics of the original dataset.

4.2.3 Differential Privacy

DP is a theoretical definition of privacy that has quickly become a standard in the privacy community [18]. First proposed by Dwork in 2006 [17], DP is popular as it provides a theoretical bound on the output data distribution. The privacy guarantee applies even in the worst-case scenario where all other entries from the original dataset have been leaked. The privacy parameters for DP are defined to quantify how much information an attacker gains when they access data released by the privacy mechanism. An assumption on attacker knowledge is unnecessary as the DP guarantee applies to any two datasets that differ by at most one element. Formally, ε-differential privacy (ε-DP) is defined as,

Definition 3 *ε-DP*
A mechanism $\mathbf{M}$ provides ε-DP if for all databases D,D' that differ in at most one element and for every $O \subseteq Range(\mathbf{M})$, we have

$$Pr[\mathbf{M}(D) \in O] \leq e^{\varepsilon} \cdot Pr[\mathbf{M}(D') \in O]$$

ε-DP applies to the mechanism $\mathbf{M}$, and not the database D or D'. A guarantee on the mechanism ensures that the formal guarantee generalizes to all possible datasets that vary by only one element. The e^{ε} term bounds the probability that an attacker can detect a difference if a given data element was or was not contained within the original dataset. A major benefit of DP mechanisms is not requiring assumptions on what information an attacker has. As a result, the privacy parameter ε is easier to interpret across specific datasets and applications when applied to the same data type.

4.3 Implementation

This section introduces two new privacy mechanisms for eye-tracking sample datasets that achieve k-anonymity and PD through generative models. First, we describe the generative models used to synthesize gaze samples during fixation and saccade events. Second, we describe the implementation of our proposed methods and the existing kalεido mechanism. Please see the Supplementary Material for pseudocode describing all three mechanisms.

4.3.1 Synthesis Models

Privacy for sample-level data is achieved by synthesizing new gaze samples (Fig. 2). The approach to gaze synthesis is first to identify fixation and saccade events and then replace gaze samples during the events with synthetic data.

Fixations Fixations are low-velocity eye movements best described as clusters of gaze positions around a fixation center. We applied a simple model that fits an anisotropic 2D Normal distribution with parameters μ_x, μ_y, σ_x, and σ_y for each fixation cluster and generated synthetic gaze samples by sampling from this distribution.

To determine the probability that a set of t gaze samples were sampled from a given 2D Normal distribution,

$$Pr\{y = \{(x_1,y_1), \cdots, (x_i,y_i), \cdots (x_t,y_t)\} \leftarrow N(\mu_x,\mu_y,\sigma_x,\sigma_y)\}, \quad (1)$$

we considered the joint probability that all of the points come from the Normal distribution $N(\mu_x,\mu_y,\sigma_x,\sigma_y)$. The joint probability for independently sampled points is computed as a product of probabilities that each point came from the same distribution $\prod_{i=1}^{t} Pr\{(x_i,y_i) = N(\mu_x,\mu_y,\sigma_x,\sigma_y)\}$. Gaze positions in this context are considered a continuous random variable defined by $N(\mu_x,\mu_y,\sigma_x,\sigma_y)$. We computed the individual probabilities by considering the cumulative distribution function (CDF) for the Normal distribution. The CDF returns probabilities that the random variable falls within a range of values a and b in the form $Pr\{(a_x,a_y) < N(\mu_x,\mu_y,\sigma_x,\sigma_y) \leq (b_x,b_y)\}$. We approximated $Pr\{(x_i,y_i) = N(\mu_x,\mu_y,\sigma_x,\sigma_y)\}$ as $Pr\{(x_i-\partial, y_i-\partial) < N(\mu_x,\mu_y,\sigma_x,\sigma_y) \leq (x_i+\partial, y_i+\partial)\}$, where $\partial = .01$ represents a sufficiently small region around the gaze position to consider. Estimating the probability from the CDF between $(x_i-\partial, y_i-\partial)$ and $(x_i+\partial, y_i+\partial)$ provides the probability that a value near (x_i,y_i) comes from the distribution $N(\mu_x,\mu_y,\sigma_x,\sigma_y)$. The probability is used to compute the Pr term in EQ. 1, thus allowing the PD Privacy Test to be applied at the fixation level.

Saccades A three parameter Gaussian function is fit to the profile of instantaneous saccade velocities computed from the raw gaze samples [25, 66]. The velocities v_i for each gaze sample (x_i,y_i) are computed as $d((x_{i-1},y_{i-1}),(x_i,y_i))$ where d is the shortest angular distance between two points on a sphere of uniform radius computed with the haversine formula [57]. The Gaussian function used to model the velocity profile of a saccade is defined as $G(a,b,c,t) = a * e^{-\frac{(t-b)^2}{c}}$, where a, b, and c control the shape of the velocity profile and $t \in [0,1]$ represents normalized saccade duration. Raw velocity values are resampled uniformly to a fixed number of values from saccade start to

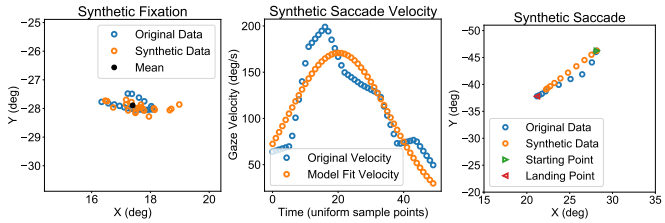

Fig. 2: Synthetic gaze data (orange circles) for fixations and saccades. Left: Samples from a fixation modeled using a 2D Normal distribution. Center: Saccade velocity profile modeled using a three-parameter Gaussian model. Right: The velocity profile is integrated to generate displacements between consecutive points, producing synthetic positions between the starting and landing point of the saccade.

end point. We fixed the number of samples to 30 for all saccade profiles. The parameters a, b, and c for each saccade are determined by applying the scipy.optimize.least_squares function to minimize the sum of squared errors for each saccade profile, $\sum_{i=1}^{30}(v_i - G(a,b,c,t_i))^2$, where $t_i = (i-1)\cdot\frac{1}{29}$.

Probabilities that an individual produced a synthetic saccade profile is required to implement the PD test. The saccade velocity values for each uniform time step i of the real dataset are used to define probability mass functions by mapping continuous values to discrete bins. Fifty bins discretized the velocity values at each time step. The bins uniformly covered a range of velocity values between 0 and 1000 degrees per second. A histogram for each individual counted the number of velocity values that fell into each bin at each time step. The counts are divided by the total number of saccades from each individual, so the sum of all probabilities are equal to one. The resulting values provide a joint probability distribution over the likelihood of producing a specific velocity value at each time step. The probabilities across time points are summed to compute the likelihood that an individual generated a synthetic saccade profile.

Synthetic saccade positions are generated from a velocity profile by computing the displacement between each sample in the event. The displacement represents the amplitude A at each time step. The amplitudes are used to generate synthetic gaze positions iteratively as $(x_{i+1}, y_{i+1}) = (x_i, y_i) + A \cdot T$, where T is a normalized vector from the saccade start point to the original landing point (Figure 2, Right).

Conditional Variational Autoencoder To achieve PD for saccades a randomized generative model is needed. The generated synthetic data must capture the characteristics of the original data without being traced to identity. Sampling saccades is more complex than fixations, which used a tractable 2D Normal distribution. The saccade profiles are represented as a discrete set of 50 velocities from saccade start to finish. Saccade direction, amplitude, duration, and individual differences affect the shape of velocity profiles, creating a multi-dimensional probability distribution. Randomly sampling the three-parameter Gaussian equation used for fitting profiles above cannot capture differences in distributions across all of these factors, and thus a deeper generative approach is necessary.

A conditional variational autoencoder (C-VAE) is a deep model that synthesizes data from distributions with user-specified conditions [62]. We explored C-VAEs as they learn from large datasets and accurately reconstruct velocity profiles given amplitude, duration, and identity labels. Prior models for PD, such as Marginals [10, 15], are limited in retaining utility as the conditions that determine the probability distributions are not used as inputs. The C-VAE model is randomized and generates a different profile each time, which is ideal for the iterative generating and testing of synthetic data for PD.

4.3.2 k-same-synth

Gaze samples from each detected fixation and saccade event are used to fit model parameters for each event in the dataset. The k-same-select sequence mechanism [15] is applied directly to the model parameters. Events are processed sequentially in the order in which they occurred. The k-anonymous model parameters are then used to sample synthetic data points for fixations and saccades.

For fixations, the μ_x, μ_y, σ_x, and σ_y parameters are processed by the mechanism to modify the centroid position of the fixation using other individuals' data and varying the spatial spread of the samples. The absolute position of the fixation within the stimulus could be shifted to a different region as a result of averaging. For saccades, the parameters of a Gaussian function model are averaged and used to construct a k-anonymous velocity profile.

4.3.3 event-synth-PD

Plausible deniability is achieved for samples by generating synthetic gaze positions for fixation and saccade events. The feature vector extracted from the events must pass the privacy criterion (Definition 2) before being released. For fixations, gaze samples are generated by randomly sampling the Gaussian distribution defined by μ, σ_x, and σ_y parameters until the privacy criteria are met. For saccades, gaze samples are generated by synthesizing new velocity profiles with the C-VAE model until the criteria are met.

We defined our own PD Event Privacy Test that determines if a synthetic fixation or saccade is k,γ-PD as an alternative to the original privacy test defined in Sec. 4.2.2. For each synthetic event, the modified privacy test loops over event parameters from other individuals. After identifying an event that passes the test for an individual, k' is incremented and the loop moves on to the next individual. The last step returns pass or fail based on whether $k' \geq k-1$. The key difference in our modified implementation is a for loop that skips to the next individual once the test has passed. The modification guarantees that instead of at least $k-1$ other data points plausibly generating the synthetic, $k-1$ *individuals* could have plausibly generated the released synthetic. Please see the Supplementary Material for pseudocode of the modified privacy test.

The difference in the guarantee achieved by the PD and PD Event privacy tests is that the k parameter refers to either data records or individuals, respectively. The original PD Privacy Test counts k' based on the number of events that satisfy the PD criterion and could provide a passing result, even though all of the records that incremented k' were from the same individual. For Event PD, k' is only incremented once per individual.

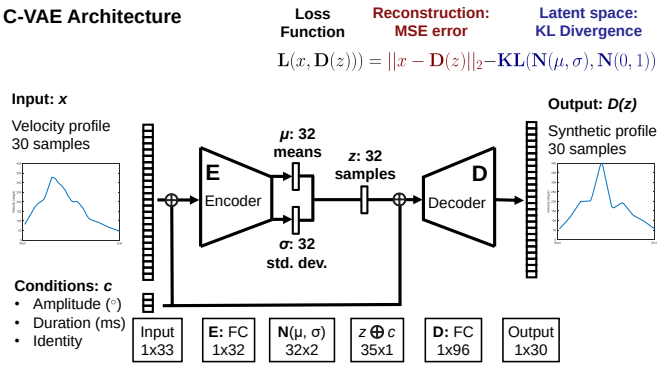

Fig. 3: C-VAE model used to generate saccade velocity profiles.

The deployed C-VAE model for saccades is used to output synthetic velocity profiles. The decoder network **D** of the C-VAE takes a randomly sampled noise vector z along with the conditions of a real saccade event, i.e., the saccade amplitude, duration, and individual identity as input, and outputs a corresponding synthetic profile. The synthetic profile captures the characteristics of the original saccade to preserve utility while also introducing random variability that will allow the extracted feature vector to pass the privacy criterion.

As shown in Figure 3, the C-VAE inputs are a velocity profile x of 30 samples concatenated with conditions c that characterize the saccade. The encoder **E** consists of a fully connected (FC) layer with 32

nodes and a ReLU activation layer. The encoder outputs 64 parameters defining a latent space of Normal distributions μ and σ. The Normal distributions are then sampled independently using inverse transform sampling to produce a noise vector z with 32 elements. The decoder $\mathbf{D}(z)$ is a one-layer FC network with 96 nodes and a linear activation layer that takes the noise vector concatenated with c, and outputs the synthetic profile. See the Supplementary Material for details on model training and optimization of parameters.

4.3.4 Kalεido

The kalεido mechanism is a composition of multiple DP definitions for processing streams of gaze samples in a real-time manner [42]. The original algorithm processed gaze samples as 2D pixel locations, while our evaluation considered 3D gaze directions. The 3D gaze directions are represented as horizontal and vertical gaze angles mapped on a unit 3D sphere.

Formally, the privacy guarantee of kalεido is defined as,

Definition 4 *(ε,w,r)-DP for gaze stream prefixes*

A mechanism $\mathbf{M} : S^g \rightarrow C^g$ *where* S^g *is the domain of all stream prefixes, satisfies (ε,w,r)-DP if for all pairs (w,r)-neighboring gaze stream prefixes* $\{S_t^g, S_t^{g'}\} \in S^g \times S^g$*, we have*

$$\forall O \in C^g, \forall t, Pr[\mathbf{M}(S_t^g) = O] \leq e^{\varepsilon} \cdot Pr[\mathbf{M}(S_t^{g'}) = O],$$

where S_t^g and S_t^g are neighboring sequences of w gaze positions prior to timestamp t, and C^g is the output set of private gaze positions.

The kalεido mechanism relies on splitting the ε DP parameter into a testing budget ε^{test} and a publishing budget ε^{pub}. The testing budget generates random noise that is added to a spatial threshold l_{thresh}. The threshold l_{thresh} plus noise acts as a fixation detector by determining if the current gaze position is close enough to the previous position to skip publishing the new position. If the distance between gaze positions is less than the threshold, then the previous gaze position is repeated in the stream.

The publishing budget determines the scale of spatial noise added to released gaze samples. The parameter h defines the ratio of testing budget to publishing budget, providing a trade-off between skipping samples more randomly and adding more spatial noise. Li et al. [42] determined l_{thresh} and h empirically and scaled them based on values of r. We set the minimum number of samples to skip t_{skip} before testing as five, l_{thresh} to one degree, and h to two in our analysis. Note that modifying parameters for the adaptive budget algorithm does not impact DP privacy, as ε does not change, but impacts utility by determining how often gaze positions are repeated.

Li et al. [42] evaluated window sizes of half a second and two seconds, and proposed a novel approach for setting the spatial bound parameter r based on ROIs within the stimulus content. We reproduced their window sizes for 100Hz eye-tracking data in our evaluation by setting w to 50 and 200 samples, respectively. For a fair comparison with k-same-synth and event-synth-PD, which do not consider stimulus content, we fixed the value of r as either the typical spatial dispersion of a fixation or the amplitude of a saccade during free viewing. Fixations are typically contained within two degrees, and the median saccade amplitude for EHTask was ten degrees during the viewing task [30].

4.4 Datasets

We evaluated the above-detailed sample privacy mechanisms on publicly available VR datasets for activity recognition and gaze prediction. The EHTask [30] dataset includes VR gaze data at 100Hz from 30 participants viewing three 360° videos. Participants viewed each video four times, performing different activities: free viewing, visual search, saliency, and tracking. A deep network was trained to classify windows of gaze and head data into the four activity classes.

The DGaze [31] dataset included VR gaze data collected at 100Hz from 43 participants that explored two 3D rendered scenes. DGaze processed saliency of scene content, tracked objects, and current gaze position to predict a future gaze position. Gaze prediction has been demonstrated in the context of foveated rendering and can help account for latency in the rendering pipeline [4, 31, 53].

4.5 Metrics

Privacy and utility metrics were computed for each dataset to determine the trade-off between the ability to re-identify users and the application of training machine-learning models.

4.5.1 Identification Rate

We computed identification rates on two trained classifiers using a Radial Basis Function network (RBFN) [15, 24, 59], with one network to classify fixation features and the other to classify saccade features. Identification rate was computed identically to related work [15]. Training and testing sets were randomly selected from unique stimuli within the dataset. Splits for training and testing were 75/25 for EHTask and 50/50 for DGaze [15]. For evaluation, one classification was made for each individual in the dataset. First, identification probabilities were computed from all of the testing data. The classification scores were then averaged within both fixation and saccade features. A final classification was made with a weighted average between the fixation and saccade scores. A weight of 0.4 was applied for fixation scores and a weight of 0.6 for saccade scores, as saccade features performed best for re-identification [15]. Identification rate was the percentage of individuals that were correctly classified with the largest class score. The reported identification rates are averaged over 20 runs of randomly selected train/test sets.

4.5.2 Activity Classification Accuracy

The EHTask classification model computes features from 1D convolutional layers applied to sequences of eye-in-head, head-in-world, and gaze-in-world samples. The output features are then fed into bidirectional GRU layers that are concatenated as input to a fully-connected network for final classification. The classification model considers data from the past ten seconds to make a prediction. Utility for EHTask was based on classifying each window of samples as the correctly. Performance was computed as accuracy $\frac{TP+TN}{TP+FP+TN+FN}$. Results were computed with 25% of the data from each task as test data by segmenting users into the train and test sets from the original paper [30]. The chance rate of guessing for EHTask is equal to 25% as there are four possible activities.

4.5.3 Gaze Prediction Accuracy

The DGaze prediction model takes eye, head, and virtual object movements as inputs processed by a 1D convolutional network combined with saliency predictions on visual content within the user's field of view. Input time windows of the past 500ms were used to train the model for predicting gaze position 100ms in the future. Utility for DGaze was measured as the angular distance between the predicted gaze position and the ground-truth future gaze position. Prediction accuracy was computed with a 60/40 train/test split from the original paper [31].

5 Results

5.1 Runtime Analysis

Privacy mechanisms were implemented in Python and processed CSV files containing raw gaze positions and labels of detected fixation/saccade events. Code was executed on a 64-bit Windows desktop with an Intel i7-6800k CPU and 16Gb of RAM. Implementation of the C-VAE models used TensorFlow (v1.13.1) with an Nvidia 1070 GPU. Reported runtimes do not measure the data pre-processing steps of importing eye-tracking data into a standard format, computing gaze velocity, removing outlier velocities, and classifying fixation or saccade events.

Table 2 presents the longest runtime across parameters for each mechanism applied to each dataset. The k-same-synth and kalεido runtimes were efficient as they processed both datasets in about two minutes or less. However, event-synth-PD took longer due to the requirement of training the C-VAE model and generating probability mass functions for saccades from the input dataset. Generating a synthetic eye-tracking dataset is feasible on a large scale for all mechanisms as the runtimes ranged between 2 and 15 minutes.

Table 2: Runtime analysis for each privacy mechanism and dataset.

Dataset	# Ppts.	Dur.	Mechanism	Runtime
EHTask	30	15 hrs	k-same-synth	2.2 mins
EHTask	30	15 hrs	event-synth-PD	14.8 mins
EHTask	30	15 hrs	kalɛido	5.2 mins
DGaze	43	5 hrs	k-same-synth	52 secs
DGaze	43	5 hrs	event-synth-PD	4.0 mins
DGaze	43	5 hrs	kalɛido	1.8 mins

Table 3: Re-identification results for the EHTask dataset. For k-same-synth, re-identification rates dropped as low as 7.5% for the largest value of k. Results from event-synth-PD dropped to 9.2% for smaller values of γ and larger values of k. Kalɛido reduced rates the most, achieving a 4% rate, which is nearly equal to chance (1/30 = 3%).

Params	RBFN identification rate % (↓)			
No mechanism	28.0%			
k-same-synth				
$k=2$	9.7%			
$k=4$	8.7%			
$k=6$	8.5%			
$k=8$	7.5%			
event-synth-PD	$\gamma=1.0$	$\gamma=1.5$	$\gamma=2.0$	$\gamma=3.0$
$k=2$	12.5%	13.5%	11.7%	13.8%
$k=4$	9.2%	12.2%	15.0%	14.2%
kalɛido	$w=50$ $r=2^\circ$	$w=50$ $r=10^\circ$	$w=200$ $r=2^\circ$	$w=200$ $r=10^\circ$
$\varepsilon=10$	10.2%	5.7%	8.0%	6.7%
$\varepsilon=5$	7.3%	4.0%	6.0%	7.2%
$\varepsilon=2$	7.7%	10.5%	9.0%	10.3%
$\varepsilon=1$	7.8%	9.3%	8.3%	6.0%

5.2 Activity Classification

Metric results for EHTask are provided in Figure 4. Visual results for each mechanism applied to EHTask can be seen in Figures 5, 6, and 7. Table 3 provides re-identification rates for several values of k. The mechanism produced re-identification rates in the range of 7.5% to 9.7%, lower than 28.0% from unmodified data. The lowest re-identification rate was 7.5% at k equals eight, which is higher than chance (1/30 = 3.3%). Rates above chance result from synthetic data affecting values of the biometric features extracted from event data. For example, the fixation features include the duration of fixations and the spatial dispersion within a fixation. The fixation synthesis method we deployed does not directly modify fixation duration, but does modify the spatial distribution that influences dispersion. Thus, any identifying trend influenced by temporal features are not guaranteed to be removed by k-same-synth.

The k-anonymous dataset was used to train an activity classification model for utility. k-anonymity introduced a loss in the activity classification utility, dropping the accuracy from 82.8% on unmodified data to as low as 29.1% (Fig. 4). A large drop in classification accuracy impacts interfaces or models that depend on recognizing the user's activity. Reasonable classification accuracy of 61.8% was achieved at k equals two, but quickly falls off and reaches chance rates at k equals eight (29.1%).

The event-synth-PD mechanism guarantees that fixation positions and the generated saccade velocity profiles are k,γ-PD. Figure 6 shows the event-synth-PD mechanism and the effect of privacy parameters k and γ on the output gaze sample positions. Synthetic horizontal and vertical gaze positions are largely unaffected for both values k, and all values of γ.

The kalɛido mechanism repeats gaze positions at high privacy levels, as the algorithm skips samples more frequently for low values of ε. Repeating gaze samples modifies the training dataset to contain less information about eye movements and prioritizes model optimization around the head movement data, which remains sampled at 100Hz. We

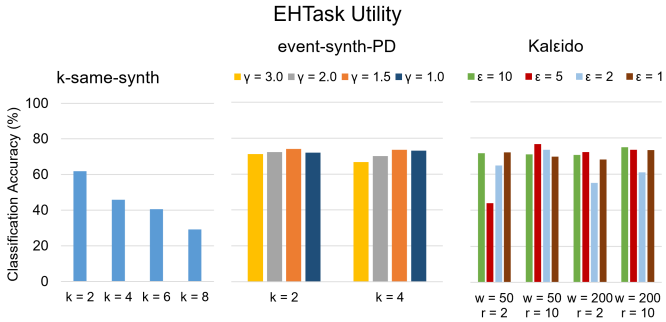

Fig. 4: Classification rates presented a downward trend for k-same-synth across values of k with rates starting near 60% and dropping to 29% with more privacy. Results from event-synth-PD presented a uniform trend near 72% for all values of k and γ. Results from kalɛido ranged from 42% to 76%; with an accuracy of 69% for parameters that achieved strong DP privacy ($r=10^\circ$, $w=200$, $\varepsilon=1$).

Table 4: Re-identification results for the DGaze dataset. Identification rates are equal to chance (1/43 = 2.3%) for unmodified data, and was lowest for k-same-synth at $k=8$.

Params	RBFN identification rate % (↓)			
No mechanism	2.3%			
k-same-synth				
$k=2$	2.0%			
$k=4$	2.3%			
$k=6$	2.1%			
$k=8$	1.1%			
event-synth-PD	$\gamma=1.0$	$\gamma=1.5$	$\gamma=2.0$	$\gamma=3.0$
$k=2$	1.2%	1.3%	1.9%	1.5%
$k=4$	1.3%	1.2%	1.5%	1.5%
kalɛido	$w=50$ $r=2^\circ$	$w=50$ $r=10^\circ$	$w=200$ $r=2^\circ$	$w=200$ $r=10^\circ$
$\varepsilon=10$	2.3%	2.3%	2.3%	2.8%
$\varepsilon=5$	2.3%	2.8%	2.1%	3.0%
$\varepsilon=2$	1.5%	2.1%	1.9%	2.3%
$\varepsilon=1$	2.6%	2.3%	4.6%	2.1%

hypothesize that the learned EHTask model classifies activity based on the sparse gaze samples and head movements relative to them, achieving comparable utility at both $\varepsilon=10$ and $\varepsilon=1$.

5.3 Gaze Prediction

Figure 8 presents results from each privacy mechanism (smaller values indicate better utility). Visual results for each mechanism applied to the DGaze prediction task can be seen in Figure 9. Table 4 provides re-identification rates for each mechanism. The DGaze dataset produced identification rates at chance with no mechanism applied. Evaluations in related work hypothesized that the low identification rates for DGaze were due to viewers only exploring two scenes in total. Viewers saw sparse environments and were instructed to follow animals moving within the scene [16]. The combination of a prescribed task and low diversity in stimuli results in features that are not reliable for user identification. This claim is supported by related work that demonstrated identification rates for free viewing were 60% higher than that of guided training sessions [41].

The k-anonymous dataset decreased utility, increasing the prediction error from 4.3° to 6.5°. Introducing 2.2° of error has a small impact on gaze prediction applications, such as foveated rendering. However, it can be compensated for with a larger foveal region parameter. For example, perceptual experiments by Guenter et al. [27] on traditional displays have found an optimal size for the foveal region between three to four visual degrees. Increasing the foveal region within this range by 2.2° to accommodate additional error would reduce the rendering speedup from a factor of ten to four.

Table 5: Summary of privacy-utility trade-offs with check marks indicating a practical trade-off for that application.

Mechanism	Guarantee	Data type	Utility	Practical trade-off
k-same-select sequence [15]	k-anonymity	Features	Document Type Classification	✓
Marginals [15]	k,γ-PD	Features	Document Type Classification	×
Exponential-DP [63]	ε-DP	Features	Document Type Classification	×
k-same-synth (ours)	k-anonymity	Samples	Activity Type Classification	×
event-synth-PD (ours)	k,γ-PD	Samples	Activity Type Classification	✓
Kalɛido [42]	ε-DP	Samples	Activity Type Classification	✓
k-same-synth (ours)	k-anonymity	Samples	Prediction	✓
event-synth-PD (ours)	k,γ-PD	Samples	Prediction	×
Kalɛido [42]	ε-DP	Samples	Prediction	×

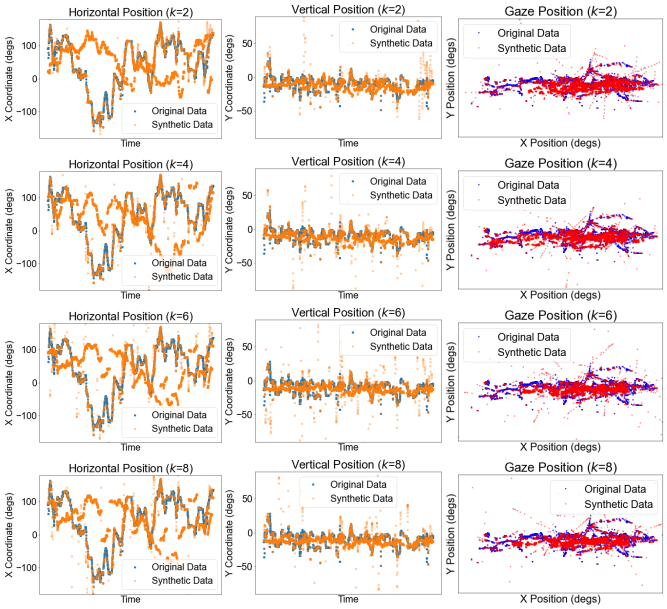

Fig. 5: Real and synthetic gaze positions for the k-same-synth mechanism from EHTask. Left Column: Horizontal position time series. Middle Column: Vertical position time series. Right Column: 2D Gaze positions in equirectangular format. Large shifts in the synthetic horizontal positions are observed as early as k equals two.

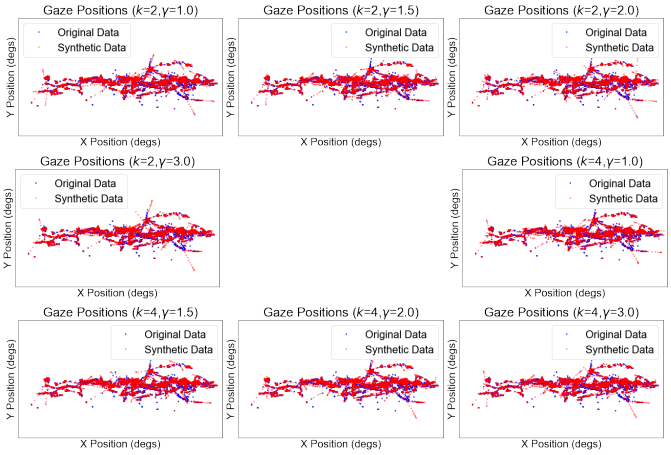

Fig. 6: Real and synthetic gaze positions for the event-synth-PD mechanism from EHTask. Synthetic positions are consistent across privacy parameters and do not vary significantly from the original data.

The k,γ-PD dataset introduced a moderate loss in utility, increasing the prediction error from the 4.3° baseline up to 9.1° across parameters. The introduced error is almost double that of the unmodified data, introducing more errors than k-same-synth.

Figure 9 shows high spatial displacement from the actual gaze position for the kalɛido mechanism. Kalɛido introduced the most error of all mechanisms at high DP privacy, but introduced reasonable errors of 5.6° or less for the smallest values of r and w. Figure 8 demonstrates a linear trend for increased prediction error within each set of DP parameters as ε decreases (higher privacy).

6 Conclusion

We presented privacy mechanisms that achieved alternative privacy guarantees to DP for eye-tracking sample data to explore RQ_1 and mitigate re-identification from gaze datasets. The presented mechanisms reduce the risk of re-identification, though not all mechanisms reduced rates to chance.

Table 5 addresses RQ_2 by providing recommendations for mechanisms that achieved practical privacy-utility trade-offs for different applications. For feature datasets used for classification tasks, past literature has recommended a mechanism that achieved k-anonymity [15]. For sample datasets used on classification models, we demonstrated practical trade-offs for both event-synth-PD and kalɛido. However, between these two mechanisms kalɛido achieved the best trade-off by sparsely sampling the gaze sample positions. The EHTask model we evaluated also takes head movements as input, and retained utility by relying less on gaze data. Visually, the gaze samples produced by event-synth-PD resembled real data much more closely (Figures 6 and 7), and would generalize better to utilities that only rely on gaze data. For sample datasets used in gaze prediction, a practical trade-off that limits introduced gaze prediction error was only achieved for the k-same-synth mechanism. In comparison to DP data from kalɛido, both privacy alternatives found practical application in one of the two tasks, and provide parameters that are simple to interpret by dataset owners.

Our proposed mechanisms protect against re-identification and are evaluated against DP, as it is considered state-of-the-art for user privacy in eye-tracking sample data. However, DP provides a guarantee that extends beyond only re-identification, also providing formal protection against sensitive inferences related to gender or age [63]. Our recommendation for which mechanism to deploy depends on the context and goal of the dataset owner, and the implications of a DP privacy trade-off are broader than the computed re-identification rates. If DP is necessary, we found a clear trend showing that higher privacy (small ε) resulted in worse performance in gaze prediction, while a clear trend was not present in activity recognition. Such a result suggests that the robust guarantee of kalɛido-DP can provide a favorable trade-off for one application, but not for others. If the scope of risks for a dataset are only focused on re-identification, then we recommend the viable alternatives of k-anonymity and PD for gaze predication and activity recognition, respectively.

6.1 Limitations

Our identification results were limited to an RBFN model, although related work explored random forest [59], SVM [49], k-NNs [11] and deep network [44] models. However, the specific model does not affect formal guarantees. Our analysis was limited to the utility model

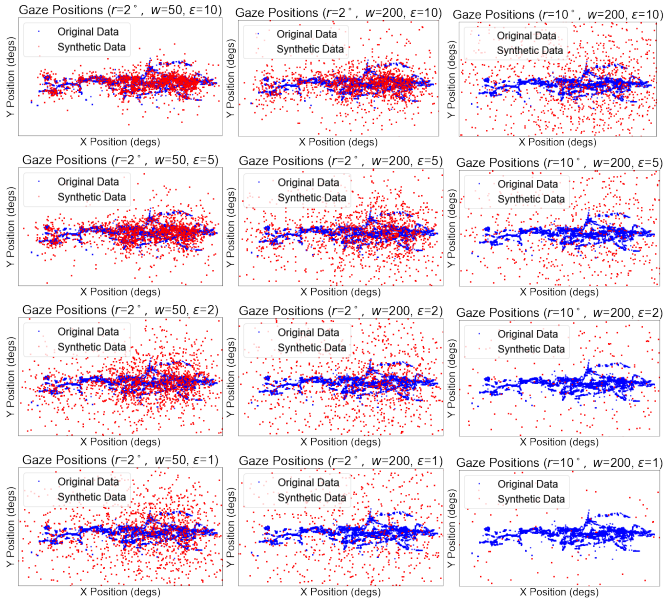

Fig. 7: Real and synthetic gaze positions for the kalɛido mechanism from EHTask. Higher levels of privacy increase spatial noise in the data and reduce temporal resolution, producing a sparse distribution of spread out gaze positions.

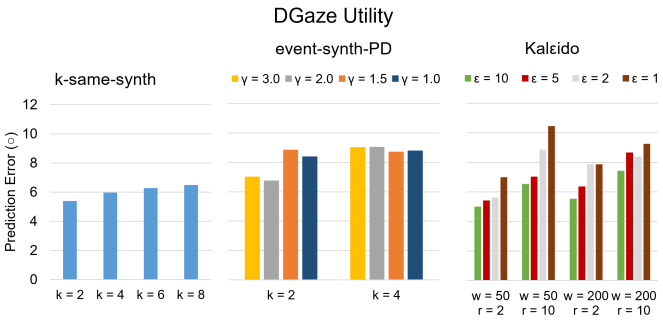

Fig. 8: Gaze prediction errors from k-same-synth increased from 5.4° to 6.5° across k, indicating a slight linear trend. Gaze errors from event-synth-PD ranged between 6.8° and 9.1° across k and γ. Gaze errors for kalɛido presented an increasing linear trend within values of r and w as ε went from low (10) to high privacy (1).

parameters reported to be optimal by the original authors. Using these parameters provided a benchmark relative to unmodified data, however, de-identified data with tweaked model parameters could result in higher utility. It would be interesting to explore trends in utility for each privacy mechanism across different hyper-parameters as well.

A fundamental limitation of privacy parameter k is the assumption that each stimulus or task has data from at least k individuals. Our results considered a classification task that included four target classes. Datasets with a larger number of classes are more difficult to classify accurately and may impact the generalization of our takeaways.

6.2 Applicability of methods to the field of VR

The presented privacy mechanisms generalize to mitigate re- identification attacks on time series data. Time series data, including eye tracking, are critical to many VR applications. Our approach to privacy-utility analysis provides a framework for exploring the protection of sensor data against re-identification attacks while still retaining utility specific to key VR applications. Relevant sensors include motion data from the head and hands of VR users, which has previously been used for accurate user identification while also being the default form of input for most VR systems [49,50]. The presented privacy mechanisms

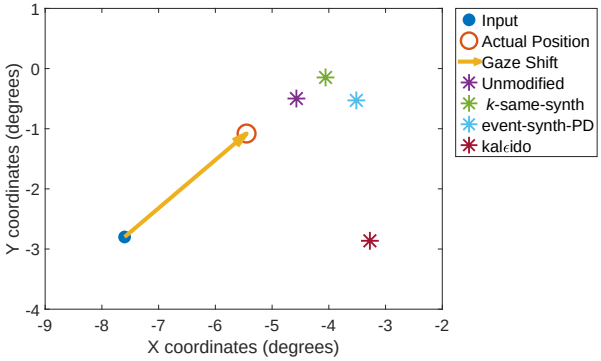

Fig. 9: Gaze predictions from DGaze trained on unmodified and private data. Colored stars indicate predictions from DGaze trained for each privacy mechanism and unmodified data. The blue dot indicates the gaze position at the time of prediction, with an arrow drawn to the actual gaze position 100ms into the future (orange circle).

for eye tracking can easily be applied to VR movement data, as they both represent 3D positional data over time.

The mechanisms can also be adapted to alternative time series data that will be integrated into VR systems, such as speech during remote collaboration, heart rate for stress detection, or brain activity for emotion recognition [8], by considering what features or representative events could be extracted from those signals. For example, with heart rate the RR interval time is a standard method for stress detection and biometric authentication [35] and provides a basis for synthesizing non-identifying data that retains utility. Once the appropriate features are identified, our proposed methods can be adapted to model bio-physical events similar to how fixations and saccades were modeled from eye movements in our work.

Privacy-preserving datasets are an important topic to explore within the VR community. As more sensors become available in the future, the ability and motivation for researchers to collect experimental data and release them publicly will also increase. While re-identification attacks on VR datasets have not yet been publicized. The research community has a unique opportunity to establish an understanding of defense methods before the risk becomes a critical issue in the field.

6.3 Future Work

Clear next steps include evaluating privacy mechanisms that consider additional sample-based VR utilities (streaming optimization [14] & adaptive interfaces [2]), alternative generative models (graph-based [10] & GANs [46]), and metric-learning biometrics [44]. The threat to privacy we considered was identifying the user from their eye-tracking data only. Attention data paired with content has the potential to violate privacy expectations concerning personalized ads, revealing biases, and identifying sexual orientation [48]. Inferences of this nature are known as biometric pyschography and indicate the emotional state or intention of a user [29]. Even de-identified data could leak mental health conditions [5] or neuro-atypicality such as autism [13]. Current privacy solutions have only explored the additional threat of gender classification with DP methods [11,22,63]. There is still a gap between the large body of work on using eye movements for sensitive characterizations in ideal lab conditions [28,61] and how frequently scenarios that produce these risks would arise in everyday use of XR.

Acknowledgments

Authors acknowledge funding from National Science Foundation (Awards FWHTF-2026540, CNS-1815883, and CNS-1562485, CNS-2206950), the National Science Foundation GRFP (Awards DGE-1315138 and DGE-1842473), and the Air Force Office of Scientific Research (Awards FA9550-19-1-0169).

REFERENCES

[1] I. Agtzidis, M. Startsev, and M. Dorr. A ground-truth data set and a classification algorithm for eye movements in 360-degree videos. *arXiv preprint arXiv:1903.06474*, 2019. 3

[2] R. Alghofaili, Y. Sawahata, H. Huang, H.-C. Wang, T. Shiratori, and L.-F. Yu. Lost in style: Gaze-driven adaptive aid for vr navigation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2019. 9

[3] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pp. 901–914, 2013. 2

[4] E. Arabadzhiyska, O. T. Tursun, K. Myszkowski, H.-P. Seidel, and P. Didyk. Saccade landing position prediction for gaze-contingent rendering. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017. 3, 6

[5] T. Armstrong and B. O. Olatunji. Eye tracking of attention in the affective disorders: A meta-analytic review and synthesis. *Clinical psychology review*, 32(8):704–723, 2012. 9

[6] A. T. Bahill and L. Stark. The trajectories of saccadic eye movements. *Scientific American*, 240(1):108–117, 1979. 3

[7] S. Berkovsky, R. Taib, I. Koprinska, E. Wang, Y. Zeng, J. Li, and S. Kleitman. Detecting personality traits using eye-tracking data. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2019. 1

[8] G. Bernal. Developing Galea: An open source tool at the intersection of VR and neuroscience. https://www.media.mit.edu/posts/galea/, 2021. Accessed: 2022-06-06. 9

[9] V. Bindschaedler and R. Shokri. Synthesizing plausible privacy-preserving location traces. In *2016 IEEE Symposium on Security and Privacy (SP)*, pp. 546–563. IEEE, 2016. 2, 4

[10] V. Bindschaedler, R. Shokri, and C. A. Gunter. Plausible deniability for privacy-preserving data synthesis. *Proceedings of the VLDB Endowment*, 10(5), 2017. 2, 4, 5, 9

[11] E. Bozkir, O. Günlü, W. Fuhl, R. F. Schaefer, and E. Kasneci. Differential privacy for eye tracking with temporal correlations. *Plos one*, 16(8):e0255979, 2021. 1, 2, 8, 9

[12] A. Bulling, J. A. Ward, H. Gellersen, and G. Tröster. Eye movement analysis for activity recognition using electrooculography. *IEEE transactions on pattern analysis and machine intelligence*, 33(4):741–753, 2010. 2

[13] K. Chawarska and F. Shic. Looking but not seeing: Atypical visual scanning and recognition of faces in 2 and 4-year-old children with autism spectrum disorder. *Journal of autism and developmental disorders*, 39(12):1663, 2009. 9

[14] S. Chen, B. Duinkharjav, X. Sun, L.-Y. Wei, S. Petrangeli, J. Echevarria, C. Silva, and Q. Sun. Instant reality: Gaze-contingent perceptual optimization for 3d virtual reality streaming. *IEEE Transactions on Visualization and Computer Graphics*, 28(5):2157–2167, 2022. 9

[15] B. David-John, K. Butler, and E. Jain. For your eyes only: Privacy-preserving eye-tracking datasets. In *ACM Symposium on Eye Tracking Research and Applications*, pp. 1–6, 2022. 2, 3, 5, 6, 8

[16] B. David-John, D. Hosfelt, K. Butler, and E. Jain. A privacy-preserving approach to streaming eye-tracking data. *IEEE Transactions on Visualization and Computer Graphics*, 2021. 1, 3, 7

[17] C. Dwork. Differential privacy. In *Automata, Languages and Programming*, pp. 1–12. Springer Berlin Heidelberg, 2006. doi: 10.1007/11787006_1 2, 4

[18] C. Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pp. 1–19. Springer, 2008. 1, 4

[19] C. Dwork, A. Roth, et al. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014. 1, 2

[20] S. Eberz, G. Lovisotto, K. B. Rasmussen, V. Lenders, and I. Martinovic. 28 blinks later: Tackling practical challenges of eye movement biometrics. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1187–1199, 2019. 1

[21] K. El Emam and F. K. Dankar. Protecting privacy using k-anonymity. *Journal of the American Medical Informatics Association*, 15(5):627–637, 2008. 3

[22] W. Fuhl, E. Bozkir, and E. Kasneci. Reinforcement learning for the privacy preservation and manipulation of eye tracking data. In *International Conference on Artificial Neural Networks*, pp. 595–607. Springer, 2021. 3, 9

[23] C. Galdi, M. Nappi, D. Riccio, and H. Wechsler. Eye movement analysis for human authentication: a critical survey. *Pattern Recognition Letters*, 84:272–283, 2016. 1, 2

[24] A. George and A. Routray. A score level fusion method for eye movement biometrics. *Pattern Recognition Letters*, 82:207–215, 2016. 1, 2, 6

[25] H. Griffith, S. Aziz, and O. Komogortsev. Prediction of oblique saccade trajectories using learned velocity profile parameter mappings. In *2020 10th Annual Computing and Communication Workshop and Conference (CCWC)*, pp. 0018–0024. IEEE, 2020. 3, 4

[26] R. Gross, E. Airoldi, B. Malin, and L. Sweeney. Integrating utility into face de-identification. In *International Workshop on Privacy Enhancing Technologies*, pp. 227–242. Springer, 2005. 2

[27] B. Guenter, M. Finch, S. Drucker, D. Tan, and J. Snyder. Foveated 3d graphics. *ACM Transactions on Graphics (TOG)*, 31(6):164, 2012. 7

[28] K. Harezlak and P. Kasprowski. Application of eye tracking in medicine: A survey, research issues and challenges. *Computerized Medical Imaging and Graphics*, 65:176–190, 2018. 9

[29] B. Heller. Watching androids dream of electric sheep: Immersive technology, biometric psychography, and the law. *Vanderbilt Journal of Entertainment & Technology Law*, 23(1):1, 2020. 9

[30] Z. Hu, A. Bulling, S. Li, and G. Wang. Ehtask: Recognizing user tasks from eye and head movements in immersive virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 2021. 6

[31] Z. Hu, S. Li, C. Zhang, K. Yi, G. Wang, and D. Manocha. Dgaze: Cnn-based gaze prediction in dynamic scenes. *IEEE transactions on visualization and computer graphics*, 26(5):1902–1911, 2020. 6

[32] B. John, S. Jörg, S. Koppal, and E. Jain. The security-utility trade-off for iris authentication and eye animation for social virtual avatars. *IEEE transactions on visualization and computer graphics*, 2020. 1

[33] B. John, S. Koppal, and E. Jain. EyeVEIL: degrading iris authentication in eye tracking headsets. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, p. 37. ACM, 2019. 1

[34] B. John, A. Liu, L. Xia, S. Koppal, and E. Jain. Let it snow: Adding pixel noise to protect the user's identity. In *ACM Symposium on Eye Tracking Research and Applications*, pp. 1–3, 2020. 1

[35] N. Karimian, D. Woodard, and D. Forte. Ecg biometric: Spoofing and countermeasures. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2(3):257–270, 2020. 9

[36] A. Karpov, J. Liberman, D. Lohr, and O. Komogortsev. Parallel oculomotor plant mathematical model for large scale eye movement simulation. *arXiv preprint arXiv:2007.09884*, 2020. 3

[37] G. Kellaris, S. Papadopoulos, X. Xiao, and D. Papadias. Differentially private event sequences over infinite streams. *Proceedings of the VLDB Endowment*, 7(12):1155–1166, 2014. 2

[38] D. Kifer and A. Machanavajjhala. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, pp. 193–204, 2011. 1

[39] O. Komogortsev, C. Holland, S. Jayarathna, and A. Karpov. 2d linear oculomotor plant mathematical model: Verification and biometric applications. *ACM Transactions on Applied Perception (TAP)*, 10(4):1–18, 2013. 3

[40] G. Lan, T. Scargill, and M. Gorlatova. Eyesyn: Psychology-inspired eye movement synthesis for gaze-based activity recognition. In *Proceedings of ACM/IEEE IPSN*, 2022. 2

[41] K. LaRubbio, J. Wright, B. David-John, A. Enqvist, and E. Jain. Who do you look like? gaze-based authentication for workers in vr. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. IEEE, 2022. 7

[42] J. Li, A. R. Chowdhury, K. Fawaz, and Y. Kim. Kalεido: Real-time privacy control for eye-tracking systems. In *29th USENIX Security Symposium (USENIX Security 20)*, 2020. 1, 2, 6, 8

[43] A. Liu, L. Xia, A. Duchowski, R. Bailey, K. Holmqvist, and E. Jain. Differential privacy for eye-tracking data. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, p. 28. ACM, 2019. 1, 2

[44] D. Lohr and O. V. Komogortsev. Eye know you too: Toward viable end-to-end eye movement biometrics for user authentication. *IEEE Transactions on Information Forensics and Security*, 2022. 1, 3, 8, 9

[45] D. J. Lohr, S. Aziz, and O. Komogortsev. Eye movement biometrics using a new dataset collected in virtual reality. In *ACM Symposium on Eye Tracking Research and Applications*, pp. 1–3, 2020. 1

[46] D. Martin, A. Serrano, A. W. Bergman, G. Wetzstein, and B. Masia.

Scangan360: A generative model of realistic scanpaths for 360 images. *IEEE Transactions on Visualization & Computer Graphics*, (01):1–1, 2022. 2, 9

[47] M. McGill. Xr and the erosion of anonymity and privacy. In *The IEEE Global Initiative on Ethics of Extended Reality Report*. IEEE, 2021. 1

[48] A. H. Mhaidli and F. Schaub. Identifying manipulative advertising techniques in xr through scenario construction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–18, 2021. 9

[49] M. R. Miller, F. Herrera, H. Jun, J. A. Landay, and J. N. Bailenson. Personal identifiability of user tracking data during observation of 360-degree vr video. *Scientific Reports*, 10(1):1–10, 2020. 8, 9

[50] R. Miller, N. K. Banerjee, and S. Banerjee. Combining real-world constraints on user behavior with deep neural networks for virtual reality (vr) biometrics. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 409–418. IEEE, 2022. 9

[51] V. Nair, G. M. Garrido, and D. Song. Going incognito in the metaverse. *arXiv preprint arXiv:2208.05604*, 2022. 1

[52] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy*, pp. 111–125. IEEE, 2008. 1

[53] A. Patney, M. Salvi, J. Kim, A. Kaplanyan, C. Wyman, N. Benty, D. Luebke, and A. Lefohn. Towards foveated rendering for gaze-tracked virtual reality. *ACM Transactions on Graphics (TOG)*, 35(6):179, 2016. 6

[54] V. Rastogi and S. Nath. Differentially private aggregation of distributed time-series with transformation and encryption. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pp. 735–746, 2010. 2

[55] D. Reilly and L. Fan. A comparative evaluation of differentially private image obfuscation. In *2021 Third IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*, pp. 80–89. IEEE, 2021. 1

[56] P. Renaud, J. L. Rouleau, L. Granger, I. Barsetti, and S. Bouchard. Measuring sexual preferences in virtual reality: A pilot study. *CyberPsychology & Behavior*, 5(1):1–9, 2002. 1

[57] C. C. Robusto. The cosine-haversine formula. *The American Mathematical Monthly*, 64(1):38–40, 1957. 4

[58] P. Samarati and L. Sweeney. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. 1998. 4

[59] C. Schröder, S. M. K. Al Zaidawi, M. H. Prinzler, S. Maneth, and G. Zachmann. 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*, pp. 1–7, 2020. 1, 6, 8

[60] R. Singel. Netflix spilled your brokeback mountain secret, lawsuit claims. *Threat Level (blog), Wired*, 2009. 1

[61] V. Skaramagkas, G. Giannakakis, E. Ktistakis, D. Manousos, I. Karatzanis, N. Tachos, E. E. Tripoliti, K. Marias, D. I. Fotiadis, and M. Tsiknakis. Review of eye tracking metrics involved in emotional and cognitive processes. *IEEE Reviews in Biomedical Engineering*, 2021. 9

[62] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28:3483–3491, 2015. 5

[63] J. Steil, I. Hagestedt, M. X. Huang, and A. Bulling. Privacy-aware eye tracking using differential privacy. In *11th ACM Symposium on Eye Tracking Research & Applications*. ACM, 2019. 1, 2, 8, 9

[64] L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05):557–570, 2002. 3

[65] R. Trimananda, H. Le, H. Cui, J. T. Ho, A. Shuba, and A. Markopoulou. Ovrseen: Auditing network traffic and privacy policies in oculus vr. *arXiv preprint arXiv:2106.05407*, 2021. 3

[66] A. Van Opstal and J. Van Gisbergen. Skewness of saccadic velocity profiles: a unifying parameter for normal and slow saccades. *Vision research*, 27(5):731–745, 1987. 3, 4

[67] A. T. Zhang and O. Le Meur. How old do you look? inferring your age from your gaze. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pp. 2660–2664. IEEE, 2018. 1

[68] Y. Zhou, T. Feng, S. Shuai, X. Li, L. Sun, and H. B.-L. Duh. Edvam: a 3d eye-tracking dataset for visual attention modeling in a virtual museum. *Frontiers of Information Technology & Electronic Engineering*, 23(1):101–112, 2022. 2