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

Brendan David-John (brendanjohn@ufl.edu), Kevin Butler, Eakta Jain

University of Florida, Computer & Information Science & Engineering, USA

Motivation: Re-identification

Re-identification of users from datasets is a privacy risk.

The Netflix Challenge public dataset was de-anonymized when paired with IMDB reviews; resulting in a woman suing Netflix as she did not want her sexual orientation revealed to her family as a result [1].

Risk: Classify identity from eye-tracking feature datasets.

Harm: Public disclosure of sensitive attributes, such as sexual orientation.

Contributions: We adapt *k*-anonymity and *k,γ*-plausible deniability for eye-tracking feature data. Our results on the MPIIDPEye dataset indicates lesser impact on utility for these mechanisms than Exponential-DP.

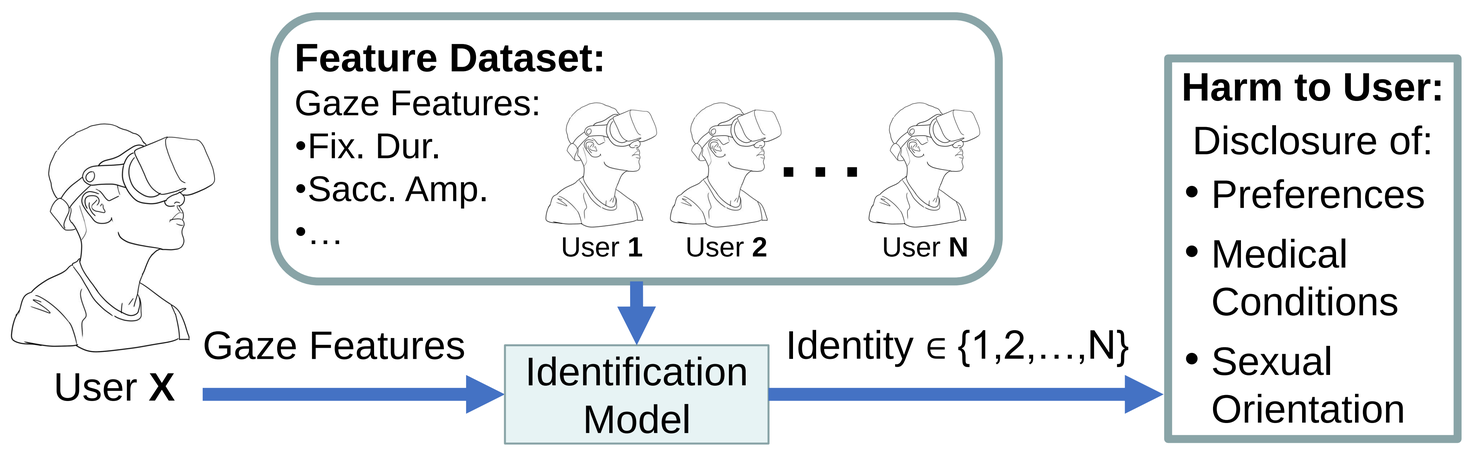

Background: Existing Privacy Definitions

***k*-anonymity**: Data from an individual cannot be distinguished from *k-1* others.

***k,γ*-plausible deniability (PD)**: Synthetic data from generative model could have plausibly been generated from *k* inputs in the real dataset.

ε-DP: A differentially private mechanism bounds the change in output probability distributions by a factor of e^{ε} when input datasets vary by only one data element.

Prior to this work, ε-DP was the only formal mechanism applied to eye tracking [2,3,4].

Characteristics \ Mechanisms	*k*-anonymity	*k,γ*-PD	ε-DP
Re-identification	✓	✓	✓
Duplicate Data	✓	✗	✗
Synthetic Data	✗	✓	✗
Noisy Data	✗	✗	✓
Impact on Utility	**Least**	**Middle**	**Most**

Adapted Mechanism 1: *k*-same-select sequence

k-same-select was originally applied to face images [5].

Randomly select *k* identities to form a group and average the sequence of feature vectors for each group.

Release *k* copies of the average feature vectors sequence for each group.

Privacy: *k* copies for each group bounds re-identification at 1/*k*.

Utility: Features are averaged within utility classes, preserving differences between classes while removing differences between identities.

Adaption: Average each groups' feature vectors across sequence.

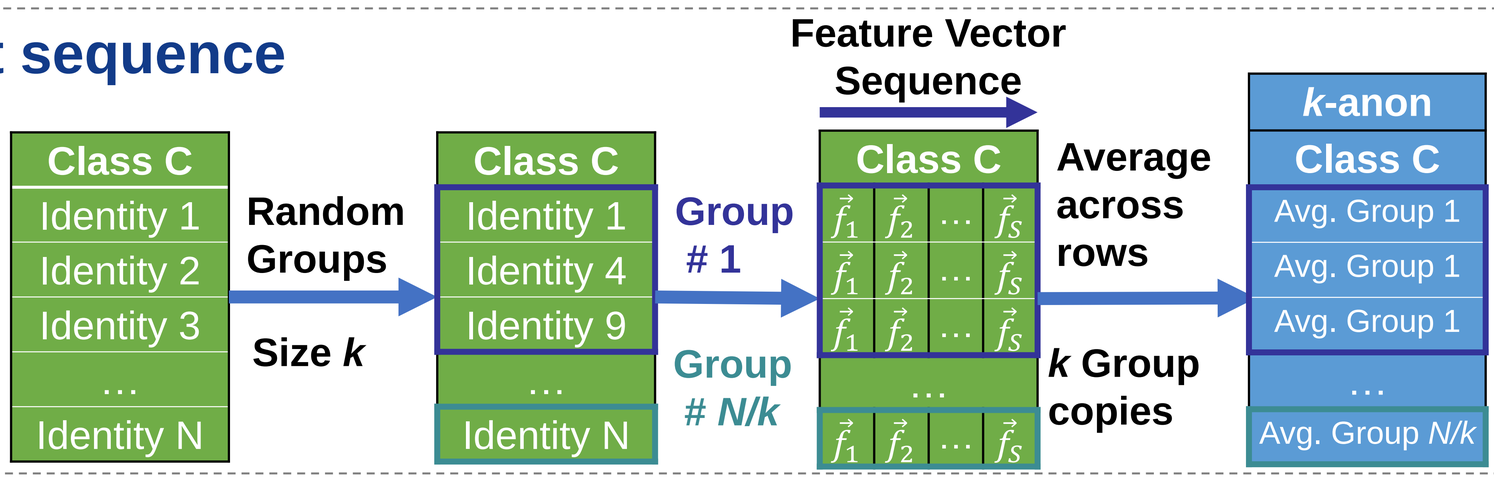

Adapted Mechanism 2: Marginals Model (PD)

The Marginals model applies to discrete data, i.e., Census data [6].

Map continuous feature values into discrete bins.

Build a Marginal distribution on the histogram of feature values and only output synthetic data that passes the PD Test (*γ* fixed at 1).

Privacy: *k* plausible seeds bounds probability of linking to real data at 1/*k*.

Utility: Synthetic data are sampled from the distribution within each class.

Adaption: Generate synthetics independently for each utility class.

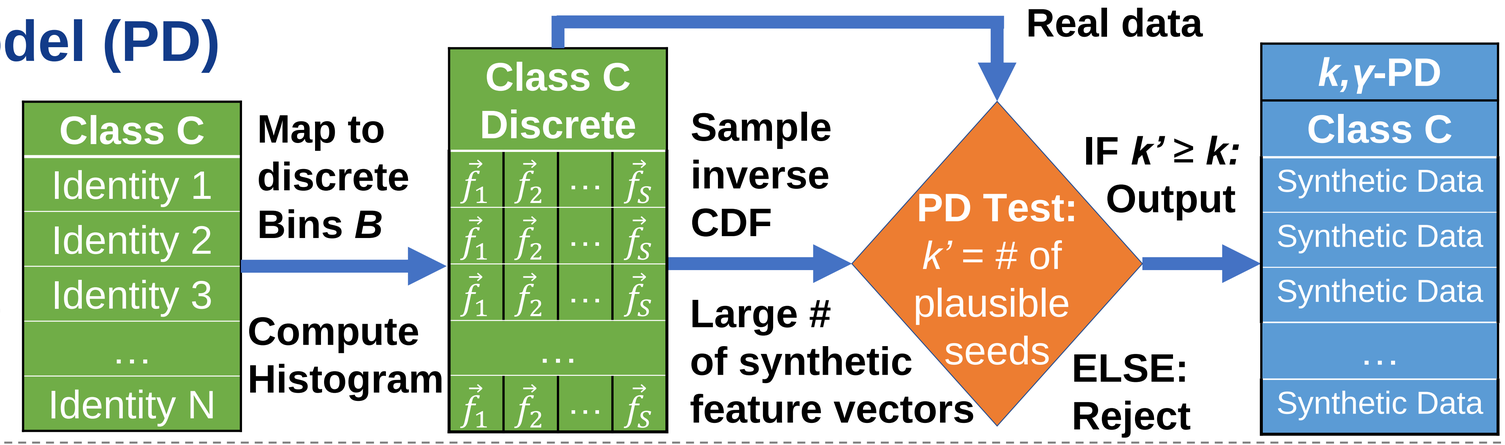

Results

Re-identification

7 VR datasets ranging from N = 13 to 130 individuals.

We computed Fixation/Saccade feature set based on [7]. (MPIIDPEye feature data used as is)

Radial Basis Function Network with 75%/25% train/test [8]. Chance identification rate is 1/N.

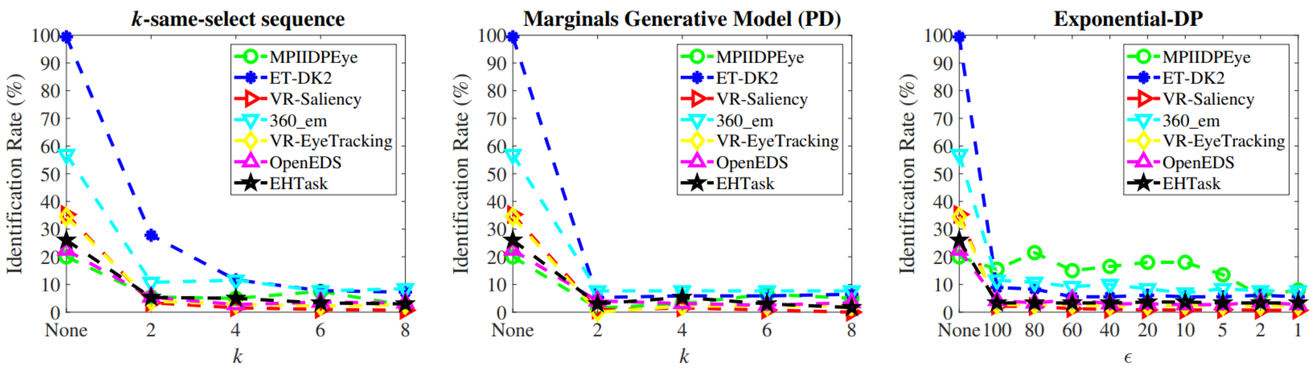

Utility

MPIIDPEye Accuracy [5]: Classify document type during reading with an SVM model.

Classes: (Comic, Newspaper, Textbook)

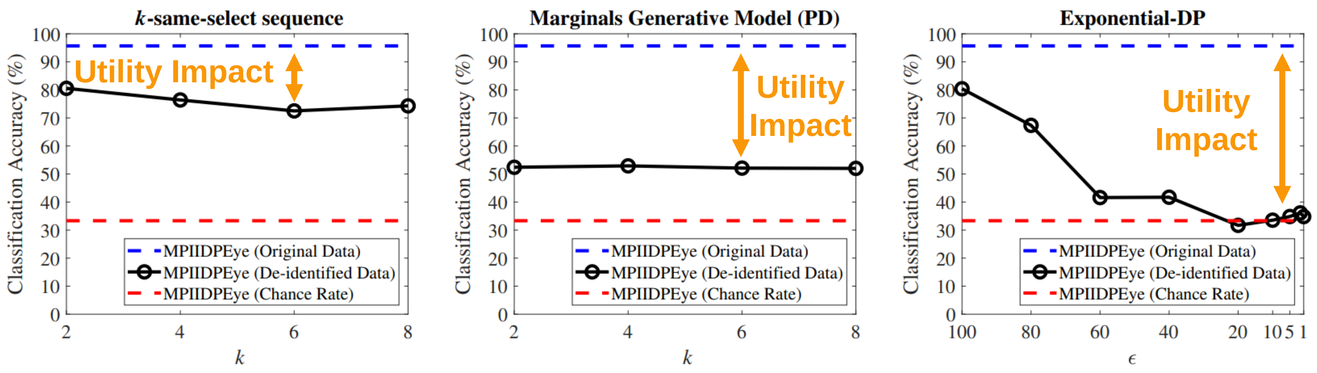

Takeaways

- Among the mechanisms we studied *k*-same-select provided the best privacy-utility trade-off at chance re-identification rates **(utility > 72%)**.
- Open-source code available: https://doi.org/10.5281/zenodo.6463850

Limitations

- Only one classification task was evaluated.
- The Marginals PD model assumes feature data columns are independent when sampling synthetic data.

Future Work

- Evaluate *k*-same-select on different classification tasks, such as intent prediction.
- Improve models for generating synthetic data with better utility for *k,γ*-plausible deniability.

References

[1] Singel R., "Netflix spilled your brokeback mountain secret, lawsuit claims", Threat Level (blog), Wired (2009).
[2] Liu, A., Xia, L., Duchowski, A., Bailey, R., Holmqvist, K., Jain, E., "Differential privacy for eye-tracking data", ETRA 2019
[3] Steil, J., Hagestedt, I., Huang, M., Bulling, A., "Privacy-Aware Eye Tracking Using Differential Privacy", ETRA 2019
[4] Bozkir, E., Günlü, O., Fuhl, W., Schaefer, R., Kasneci, E., "Differential privacy for eye tracking with temporal correlations", Plos One Vol. 16, Num. 8, 2021
[5] Gross, R., Airoldi, E., Malin, B., Sweeney, L. "Integrating utility into face de-identification", PET 2005
[6] Bindschaedler, V., Shokri, R., Gunter, C., "Plausible Deniability for Privacy-Preserving Data Synthesis", VLDB 2017
[7] George, A., Routray, A., "A score level fusion method for eye movement biometrics", Pattern Recognition Letters 2016
[8] David-John, B., Hosfelt, D., Butler, K., Jain, E., "A privacy-preserving approach to streaming eye-tracking data", IEEE TVCG 2021

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.

UF UNIVERSITY of FLORIDA

UF Herbert Wertheim College of Engineering
Department of Computer & Information Science & Engineering
UNIVERSITY of FLORIDA