Let It Snow: Adding pixel noise to protect the user's identity

Brendan John
University of Florida
brendanjohn@ufl.edu

Ao Liu
Rensselaer Polytechnic Institute
liua6@rpi.edu

Lirong Xia
Rensselaer Polytechnic Institute
xial@cs.rpi.edu

Sanjeev Koppal
University of Florida
sjkoppal@ece.ufl.edu

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

ABSTRACT

Optical eye trackers record images of the eye to estimate gaze direction. These images contain the iris of the user. While useful for authentication, these images can be used for a spoofing attack if stolen. We propose to use pixel noise to break the iris signature while retaining gaze estimation. In this paper, we present an algorithm to add "snow" to the eye image and evaluate the privacy-utility tradeoff for the choice of noise parameter.

CCS CONCEPTS

• **Security and privacy → Privacy protections**.

KEYWORDS

Privacy, Eye Tracking, Iris Recognition

ACM Reference Format:
Brendan John, Ao Liu, Lirong Xia, Sanjeev Koppal, and Eakta Jain. 2018. Let It Snow: Adding pixel noise to protect the user's identity. In *Proceedings of 1st International Workshop on Privacy and Ethics in Eye Tracking (PrEThics), June 02, 2020, Stuttgart, Germany*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION AND THREAT MODEL

Eye tracking is expected to be integrated in the next generation of automobiles, educational hardware, and virtual and mixed reality devices. Optical eye trackers typically image the eye in order to process the pixel data into an estimate of the user's gaze direction. Previous work has highlighted how these images contain the iris biometric of the user and may be used for iris authentication if stolen, and has proposed defocus to secure the iris biometric [4, 5]. We argue that in addition to a security risk, the inclusion of the iris biometric also poses a threat to privacy: there is no notion analogous to "private browsing" if the platform can always identify the user biometrically. This paper presents preliminary experiments on a novel pixel noise privacy mechanism to remove the iris signature from the eye image prior to using it for gaze estimation.

Our first contribution is a novel privacy mechanism, "snow", that can be easily applied to eye images. Our second contribution is an empirical evaluation of the privacy-utility trade-off of this mechanism. Our third contribution is a proposal for integrating the mechanism in hardware. Our fourth contribution is a description of a differentially private random mechanism for pixel noise.

PrEThics'20, June 02, 2020, Stuttgart, Germany
2018. ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

2 RELATED WORK

In response to privacy considerations raised by researchers in the eye tracking and pervasive computing communities [3, 7], and a survey by Steil and colleagues [10], there is active research on identifying threat models and proposing solutions [1, 4, 5, 8, 10, 11]. This work has focused on the threat to private information from the scene camera, for example, when entering a PIN number at a bank ATM [11], created differentially private saliency maps to prevent an adversary from obtaining a particular individual's point of regard from a publicly released aggregate saliency map [8], and added noise to the fixation locations of users to prevent user recognition and gender identification while retaining the features needed to classify the type of document being read [10]. The closest work to ours is by John and colleagues [4, 5] where the authors remove the iris signature from the eye images by defocusing it prior to gaze estimation. They evaluated the impact of changing the blur kernel on pupil detection rate, gaze accuracy, and the perceived attributes of a virtual avatar animated using this gaze data. Our work is related in that we seek to remove the iris signature from the eye images, however, the noise mechanism we propose is novel.

Our "snow" mechanism is closest to privacy methods proposed by Pittaluga and colleagues for thermal cameras that exploit exposure, gain, and bias voltages at the hardware level to block personal information while permitting object tracking and gesture recognition [9]. We are motivated by this work to create a hardware integration of the "snow" mechanism for future- eye trackers.

3 PROPOSED PRIVACY MECHANISM (SNOW)

We propose the addition of pixel-level noise to the eye image prior to gaze estimation. The mechanism is derived from salt-and-pepper noise, or "snow", wherein pixels randomly saturate. This noise is jarring to viewers of an image or video, but, as we show in this paper, has unexpected benefits.

Let us denote a grayscale image as $I(x)$ where x refers to the index of each pixel in the image. For example, for a $10x10$ image, $x = [1, 2, ...100]$. The intensity of the pixel x is represented by $I(x)$, where $I(x) \in [0, 255]$. A subset $\mathcal{S}$ of size $p \cdot I_{rows} \cdot I_{cols}$ indices are randomly selected, and a new image $I'(x)$ is created such that

$$I'(x) = \begin{cases} 127 & x \in \mathcal{S} \\ I(x) & x \notin \mathcal{S}. \end{cases} \quad (1)$$

Intuitively, as p increases, a greater percentage of the eye image contains noisy content. We find that pixel-level noise does not impact pupil or glint tracking, as these features are tracked as a 'blob' of bright or dark pixels that can still be extracted after "snow"

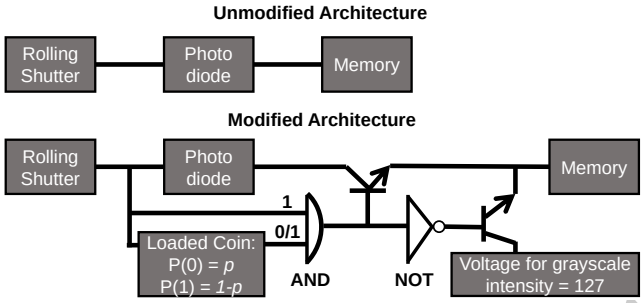

Figure 1: Modified circuit diagram for "snow".

is applied. This is an advantage over blur-based mechanisms [4, 5], as creating blur reduces pupil detection rate.

4 EMPIRICAL RESULTS

We use a dataset of eye images collected by John et al. [4]. This dataset contains 15 individuals who were eye tracked while viewing a sequence of five fixation targets presented on a computer monitor for four seconds each. Eye images were recorded using Pupil Labs software and the Pupil Pro eye tracker [6]. Before and after target viewing, eye images were also recorded for each individual for the purpose of serving as a reference set for iris based authentication. Iris authentication was performed using IrisSeg [2] and an open source implementation of iris codes[1]. The eye images were input to our noise mechanism, which created private eye images based on the "snow" mechanism and the parameter p. The private eye images were input to Pupil Labs software for pupil detection and gaze estimation.

Table 1 presents the Correct Recognition Rate (CRR) which is the percentage of private eye images that matched the reference for the individual. The gaze error is the average distance in degrees of visual angle between estimated points of regard and the associated fixation target locations. The third column presents the pupil detection rate, which is computed as the percentage of private eye images during target viewing where the pupil confidence score reported by the software was greater than 80%. Our evaluation indicates that setting 10% of the pixels to "snow" reduces the correct recognition rate to less than 20% with a negligible change in gaze estimation accuracy.

5 PROPOSED HARDWARE INTEGRATION

The "snow" mechanism can be integrated into the imaging system of a commodity eye tracker. We present a prototype block diagram where the modulation of a subset of pixels for each captured image is accomplished in hardware by using a Loaded Coin-flip to determine which pixels will be "snow". This mechanism is illustrated as a flow diagram in Figure 1. A random probability based on p is used to determine if the pixel value should be output, or a value of 127.

6 DIFFERENTIAL PRIVACY TO BENCHMARK PRIVACY

Differential privacy (DP) is a widely accepted notion of privacy. If a user uploads her data to an eye-tracking data bank in a differentially private manner, the adversary will not be able to recover personal information by comparing with other information from the user[2]. Next, we formally define DP, which can be used to benchmark the privacy of different eye-tracking algorithms.

[1]https://www.peterkovesi.com/studentprojects/libor/

p	CRR (%)		Gaze Error (°)		Samples > 0.8 (%)	
	μ	σ	μ	σ	μ	σ
0.00	78.6	14.9	1.40	0.40	76	14
0.01	41.7	24.7	1.40	0.40	71	18
0.05	29.4	17.6	1.54	0.44	71	17
0.10	17.8	14.3	1.51	0.46	71	18
0.15	12.6	13.6	2.10	0.81	70	19
0.20	9.8	12.6	2.10	1.01	69	17
0.25	7.9	11.4	2.43	0.96	69	14
0.30	6.4	11.2	2.45	1.06	69	13
0.35	5.1	9.8	2.39	0.90	68	14
0.40	4.3	9.1	2.27	0.93	64	19
0.45	3.8	8.5	2.30	0.77	63	17
0.50	3.1	7.4	2.36	0.69	63	17

Table 1: Privacy-utility trade-off for different values of p. Privacy is defined as a reduction in correct recognition rate. Utility is defined as accuracy of gaze estimation.

Definition 6.1 (DP). Let s denote a randomized algorithm and $\mathcal{S}$ be any subset of the image space of s. Then, we say s is $(\epsilon, \delta)-$differentially private if for any $\mathcal{S}$ and any pair of neighboring inputs $\boldsymbol{x}$ and $\boldsymbol{x'}$,

$$Pr\left[s(\boldsymbol{x}) \in \mathcal{S}\right] \leq e^{\epsilon} Pr\left[s(\boldsymbol{x'}) \in \mathcal{S}\right] + \delta \quad (2)$$

In (2), the probability comes from the randomness in s. Smaller ϵ and δ corresponds to stronger privacy guarantee. For the problem studied in this paper, algorithm s take images as inputs and output gaze directions. The randomized saturation process of "snow", defined in Section 3, provides randomness to s. In the following theorem, we show the DP property of the "snow" mechanism. We also note that any algorithms using "snow" as a first step will have no worse privacy than "snow" according to the post-processing property of DP.

Theorem 6.2. *"snow" with $p = 1-\delta$ is $(0, \delta)$-differentially private.*

Proof. We first take out one pixel x from the image $\boldsymbol{x}$. By the definition of snow, for any $x \neq 127$,

$$\Pr[snow(x) = x] = 1 - p \quad \text{and} \quad \Pr[snow(x) = 127] = p.$$

Then we study the difference of probability of different outputs. For any $x' \neq x$: $\Pr[snow(x) = x] - \Pr[snow(x') = x] = 1 - p$.[3]
Similarly, $\Pr[snow(x') = x'] - \Pr[snow(x) = x'] = 1 - p$.
For any output x'' such that $x'' \neq x$ and $x'' \neq x'$, we have,
$\Pr[snow(x) = x''] - \Pr[snow(x') = x''] = 0$
Because the noise of other pixels are independent from the selected one. Theorem 6.2 follows by the definition of DP. □

7 CONCLUSION AND FUTURE WORK

In this work, we presented and evaluated an additive noise privacy mechanism which we call "snow" as it creates the same type of pixel noise that was found in old televisions. We also discuss

[2]We assume the user also uploads other information in a differentially private manner.
[3]Here, x and x' corresponds to the pixel input and output from the data bank respectively. We let x, x' and x'' also denote the value of pixel.

an application of DP to this randomized algorithm. Our preliminary experiments indicate that we can replace up to 50% of the pixels in the eye image with snow and have less than 2.5° error in gaze estimation. Because the amount of acceptable error in gaze estimates depends on the intended application, future work might conduct user studies to determine people's preferences for privacy versus utility for different gaze-based applications. In future work, we plan to assess performance when using state-of-the-art deep learning based models for iris authentication. Another interesting direction would be to prototype the proposed hardware integration for various eye tracker designs.

REFERENCES

[1] Efe Bozkir, Ali Burak Ünal, Mete Akgün, Enkelejda Kasneci, and Nico Pfeifer. 2019. Privacy Preserving Gaze Estimation using Synthetic Images via a Randomized Encoding Based Framework. *arXiv preprint arXiv:1911.07936* (2019).

[2] Abhishek Gangwar, Akanksha Joshi, Ashutosh Singh, Fernando Alonso-Fernandez, and Josef Bigun. 2016. IrisSeg: A fast and robust iris segmentation framework for non-ideal iris images. In *ICB 2016*. IEEE, 1–8.

[3] Eakta Jain. 2018. Who Watches the Watchmen: Eye Tracking in XR. *Dagstuhl Seminar 18252* (2018).

[4] B John, S Jorg, S Koppal, and E Jain. 2020. The Security-Utility Trade-off for Iris Authentication and Eye Animation for Social Virtual Avatars. *IEEE transactions on visualization and computer graphics* (2020).

[5] Brendan John, Sanjeev Koppal, and Eakta Jain. 2019. EyeVEIL: degrading iris authentication in eye tracking headsets. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. ACM, 37.

[6] Moritz Kassner, William Patera, and Andreas Bulling. 2014. Pupil: an open source platform for pervasive eye tracking and mobile gaze-based interaction. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing: Adjunct publication*. ACM, 1151–1160.

[7] Daniel J Liebling and Sören Preibusch. 2014. Privacy considerations for a pervasive eye tracking world. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. ACM, 1169–1177.

[8] 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.

[9] Francesco Pittaluga, Aleksandar Zivkovic, and Sanjeev J Koppal. 2016. Sensor-level privacy for thermal cameras. In *2016 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–12.

[10] Julian Steil, Inken Hagestedt, Michael Xuelin Huang, and Andreas Bulling. 2019. Privacy-Aware Eye Tracking Using Differential Privacy. In *11th ACM Symposium on Eye Tracking Research & Applications*. ACM.

[11] Julian Steil, Marion Koelle, Wilko Heuten, Susanne Boll, and Andreas Bulling. 2019. PrivacEye: privacy-preserving head-mounted eye tracking using egocentric scene image and eye movement features. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*. ACM, 26.