Jingjie Li, Amrita Roy Chowdhury, Kassem Fawaz, Younghyun Kim
University of Wisconsin–Madison

KalεIDO: Real-Time Privacy Control for Eye-Tracking Systems

30th USENIX Security Symposium

MADS&P
Security and Privacy Research Group
at UW-Madison

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Eye-Tracking, an Emerging Human-Computer Interface

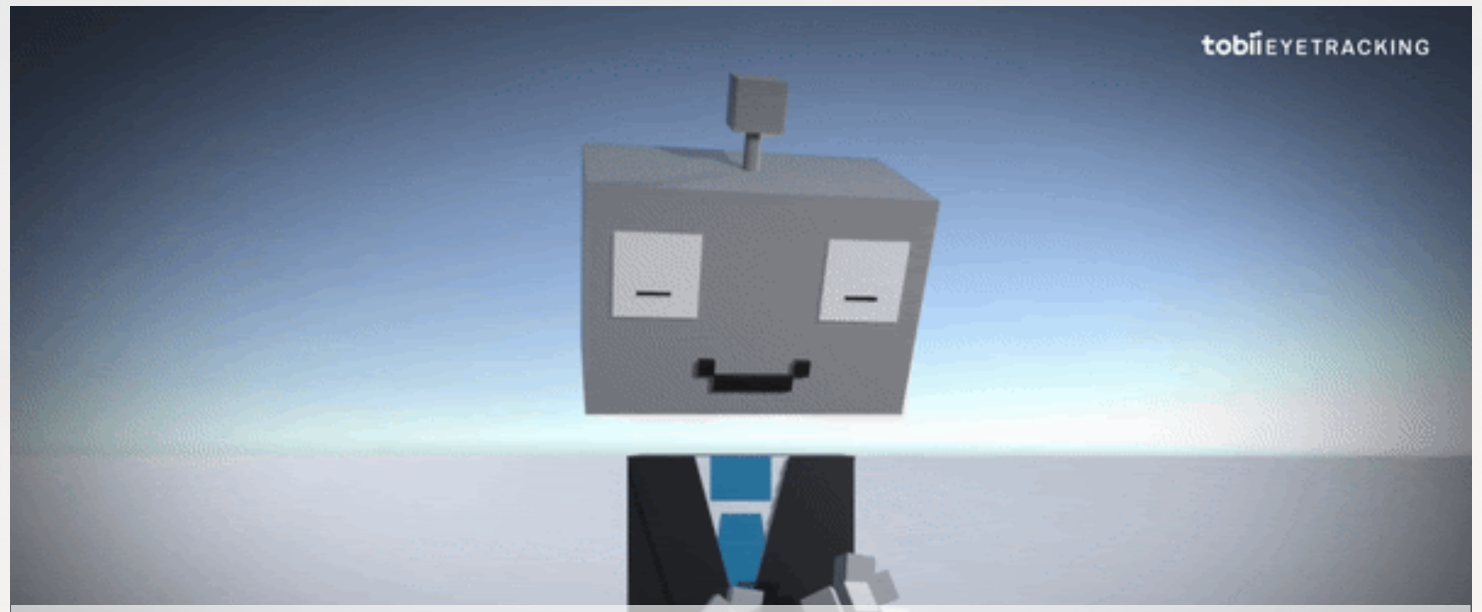

Social avatar

Foveated rendering

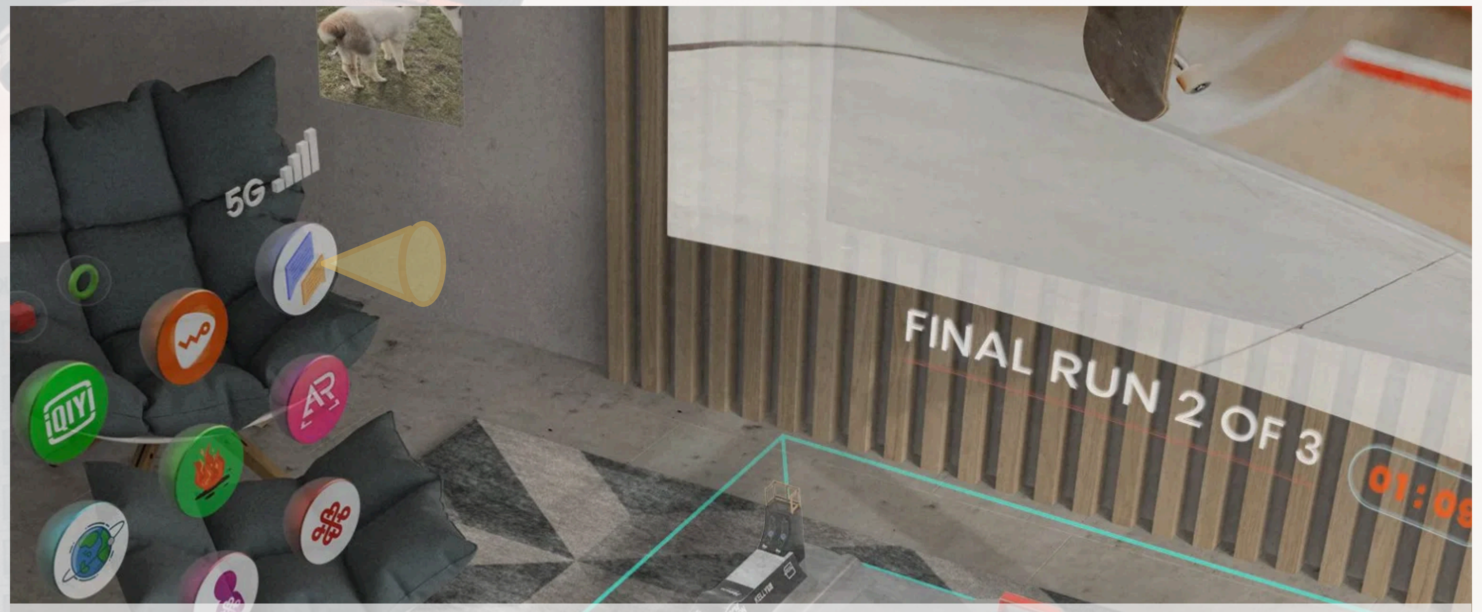

Event triggering

- Eye gazes continuously tracked by cameras
- Enables hands-free interaction
- Pervasively equipped in mixed reality

Background on Eye-Tracking Data

- **Eye gaze data:** a streaming data of timestamped location tuples (x,y,t)
- **ROI** on the visual scene attracts eye gazes
- **Fixation:** a cluster of concentrated eye gazes
- **Saccade:** gazes traveling rapidly from one fixation to another

Privacy Threat on Eye-Tracking Data

- Spatial distribution of absolute gaze positions
- Aggregate statistics of distribution over time

Leaking psycho/physiological traits

- **Psychological:** implicit interest, cultural background, personality traits, etc.
- **Physiological:** health condition (Alzheimer's, vision condition), biometric identity, etc.

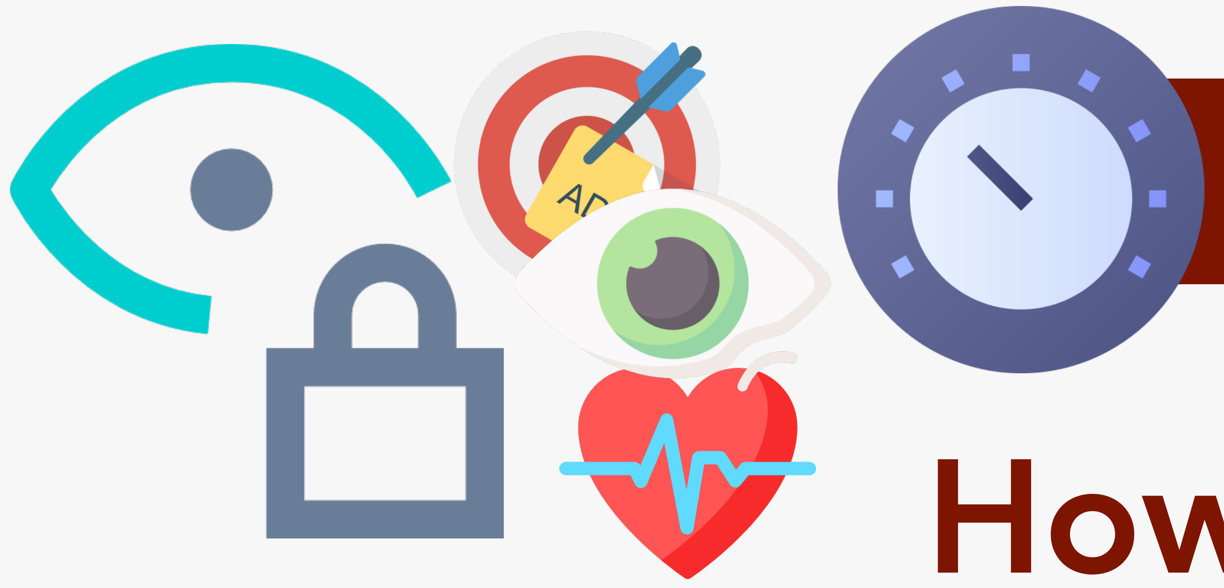

PRIVACY

Kalεido's research question

How can we control the privacy while preserving real-time utilities of eye tracking?

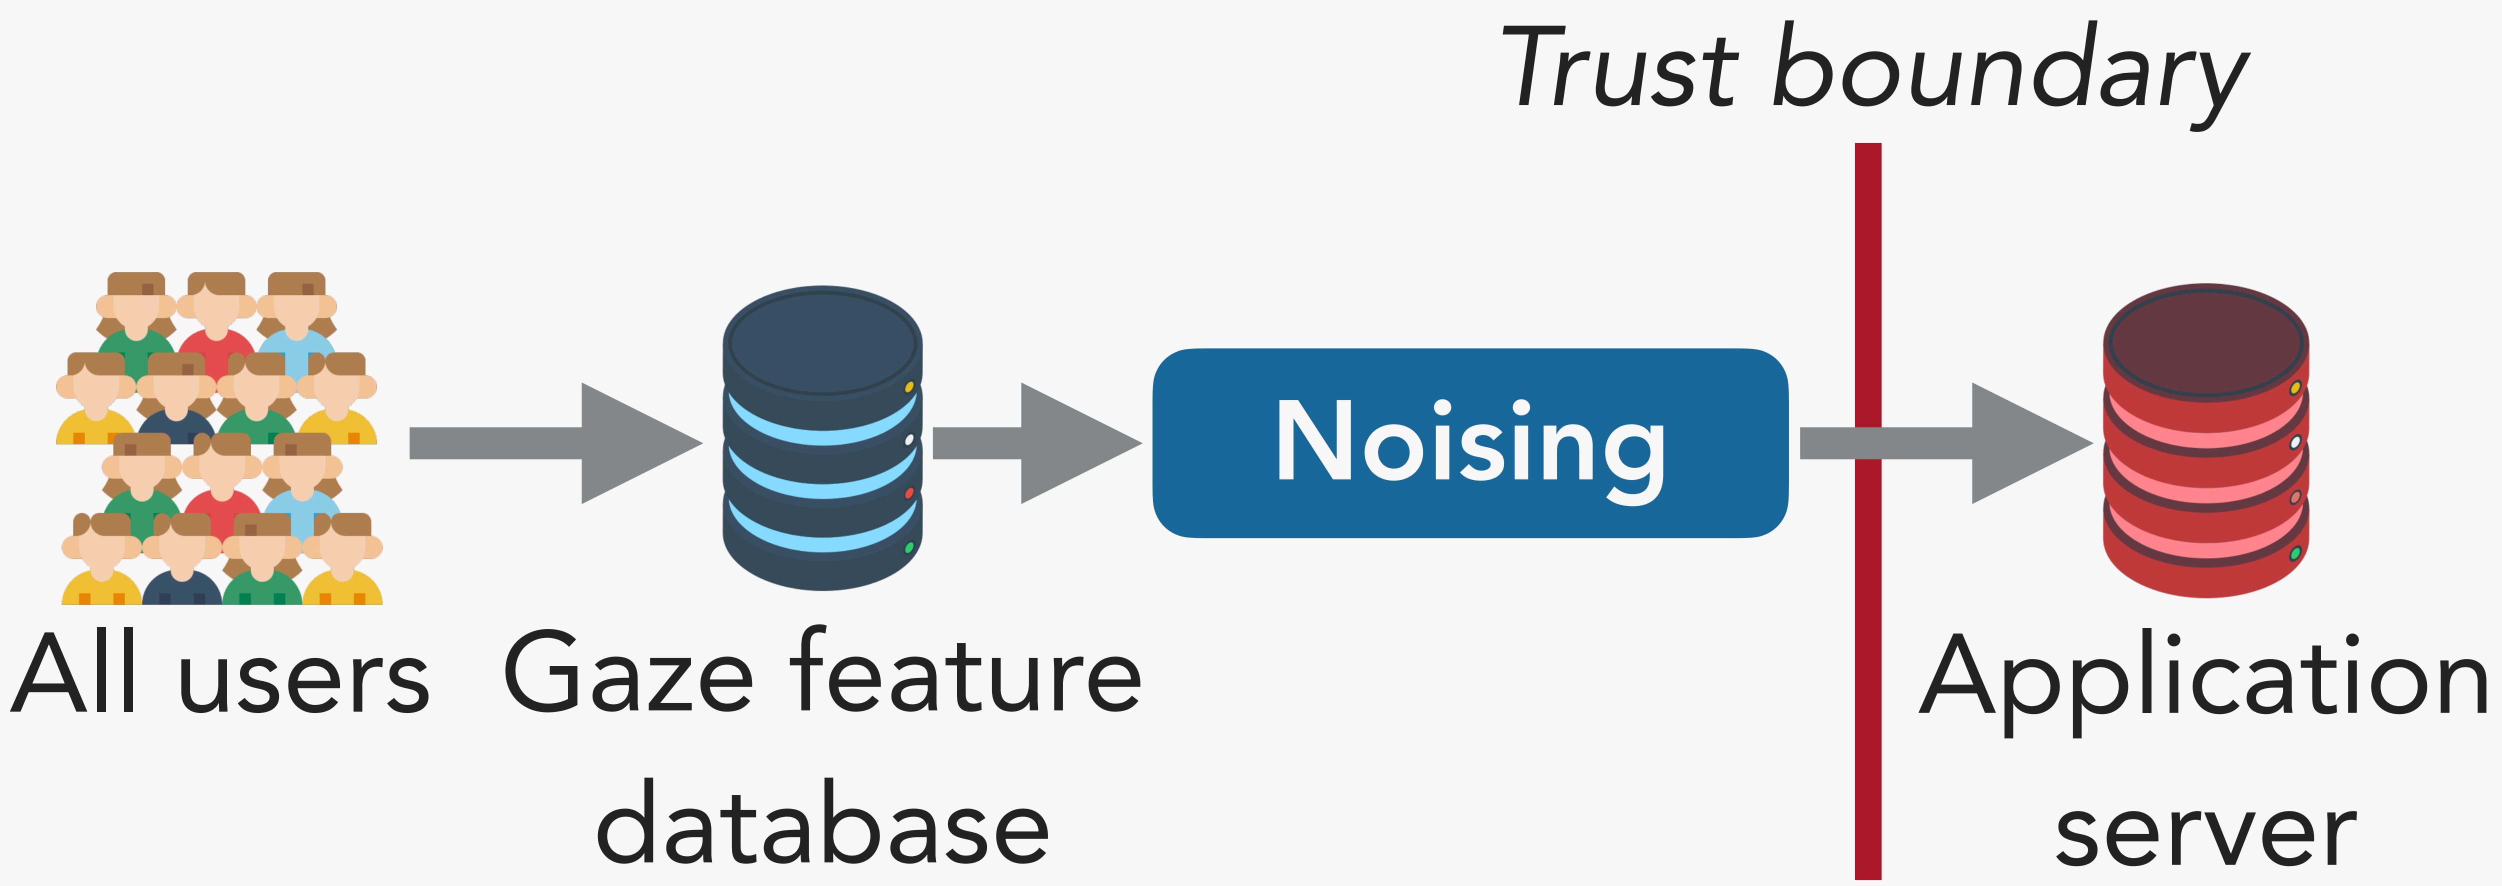

- Existing designs provide no formal guarantee (Hagestedt et al. 2020) or only allow offline release (Steil et al. 2019)
- Not suitable for real-time apps

KALεIDO: OVERVIEW

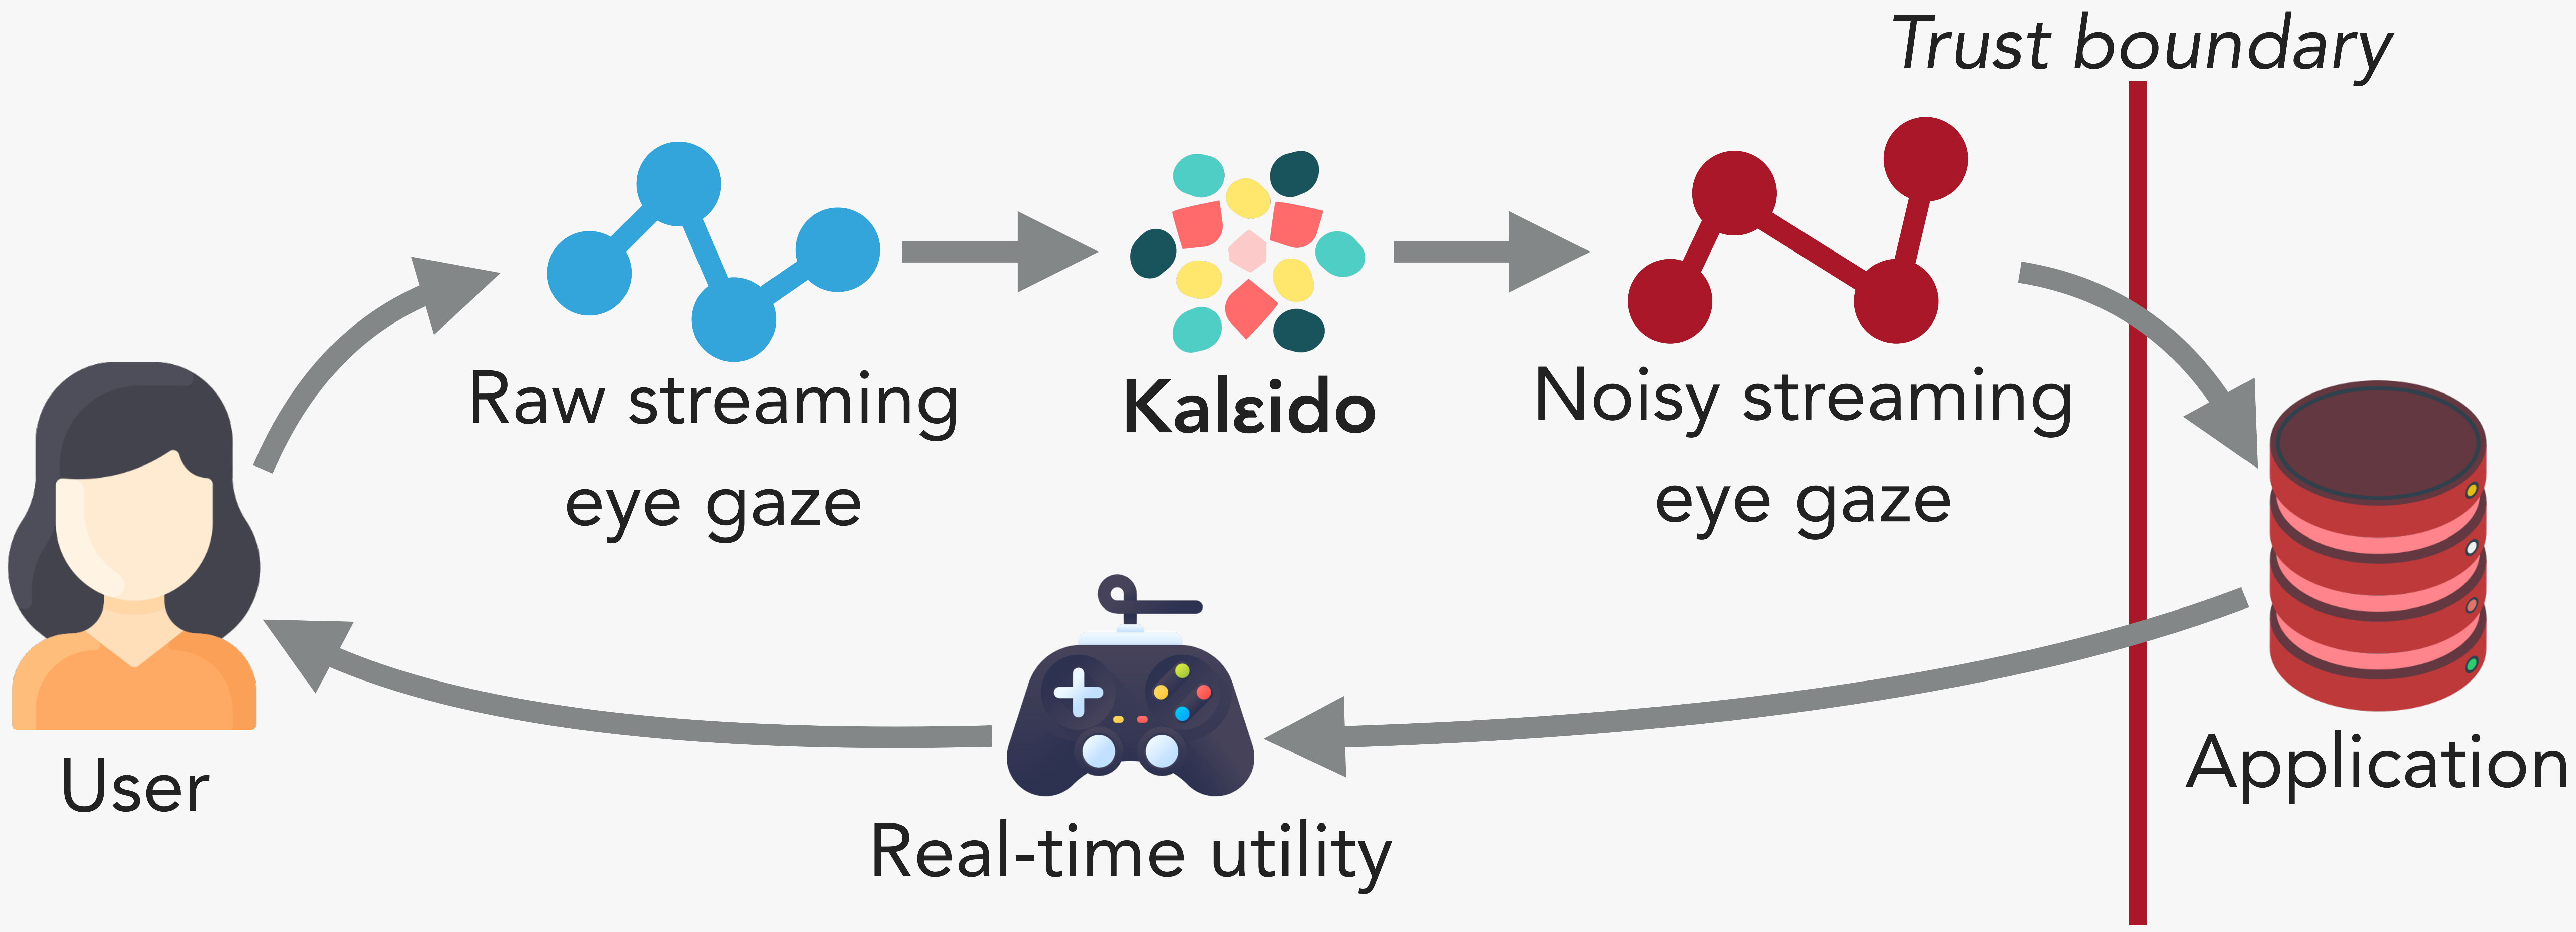

- **Formal privacy guarantee** on eye gaze streams by local differential privacy (LDP)
- **Seamless integration** with real-time eye-tracking ecosystems
- **Ease of use** by automated privacy configuration

KALεIDO: PRIVACY DEFINITION

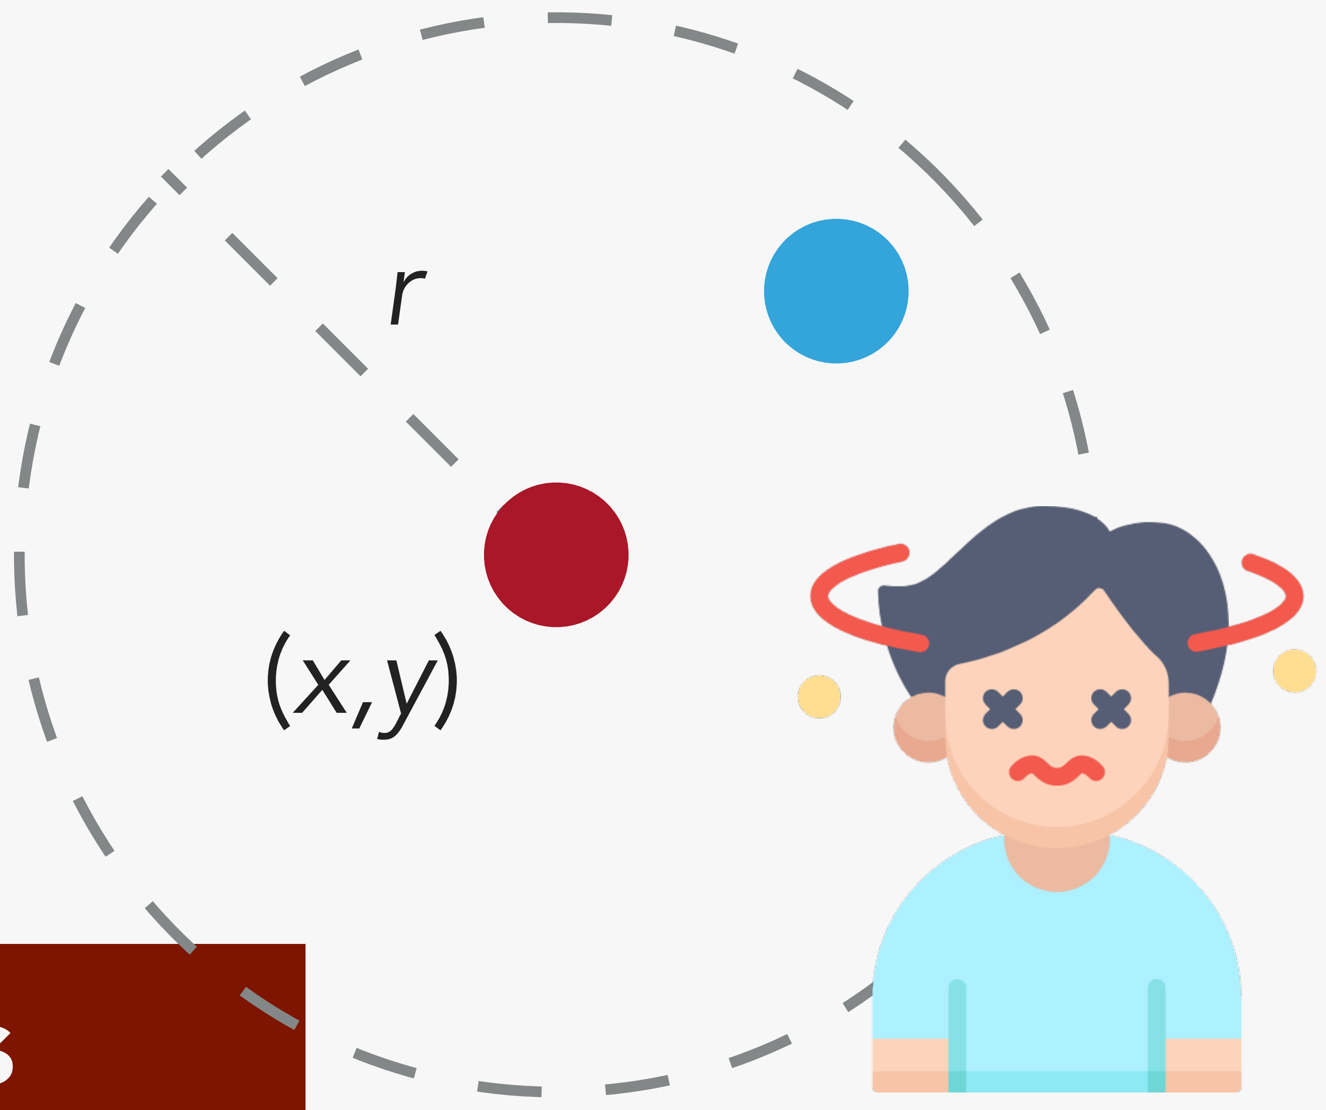

Privacy of gaze positions

Spatial information of eye gazes: primary source of sensitive information

(**ε**,*r*)-geo-indistinguishability (Andrés et al. 2013) noising $\mathscr{M} : \mathscr{G} \mapsto \mathscr{Z}$ ensures that for all pairs of inputs $(g, g') \in \mathscr{G} \times \mathscr{G}$ such that $d(g, g') \leq r$, $\forall S \subset \mathscr{Z}, Pr[\mathscr{M}(g) \in S] \leq e^{\epsilon} Pr[\mathscr{M}(g') \in S]$

KALεIDO: PRIVACY DEFINITION

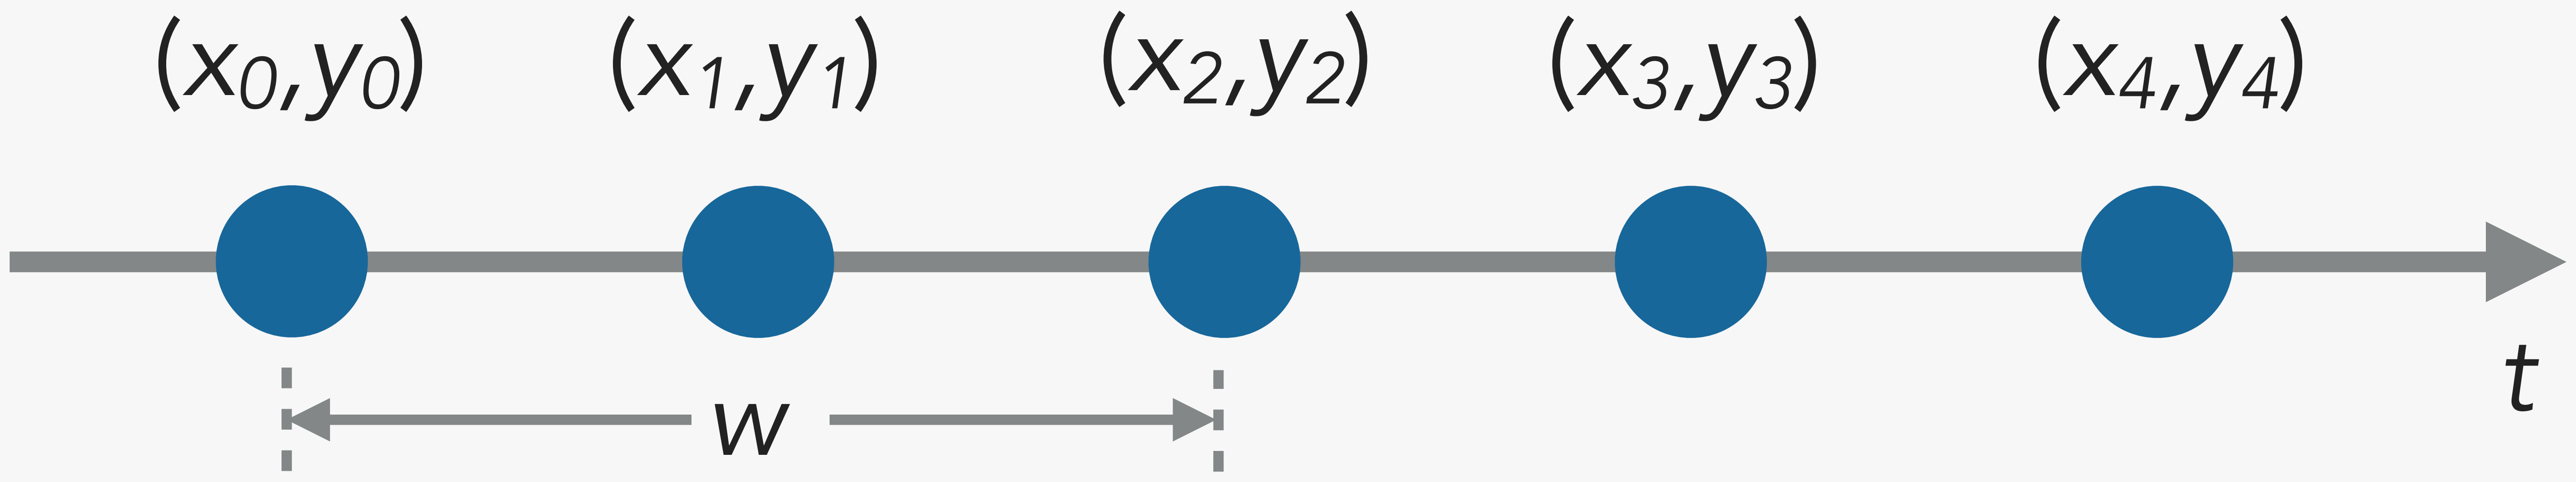

Privacy for gaze streams

Real-time streaming data: realistic format for eye-tracking interaction

(**ε**,*w*,*r*)-geo-ind. for gaze streams by leveraging *w*-event privacy (Kellaris et al. 2014) to protect the spatial distribution of any gaze trajectory formed over **any window of duration *w***

KALεIDO: IMPLEMENTATION

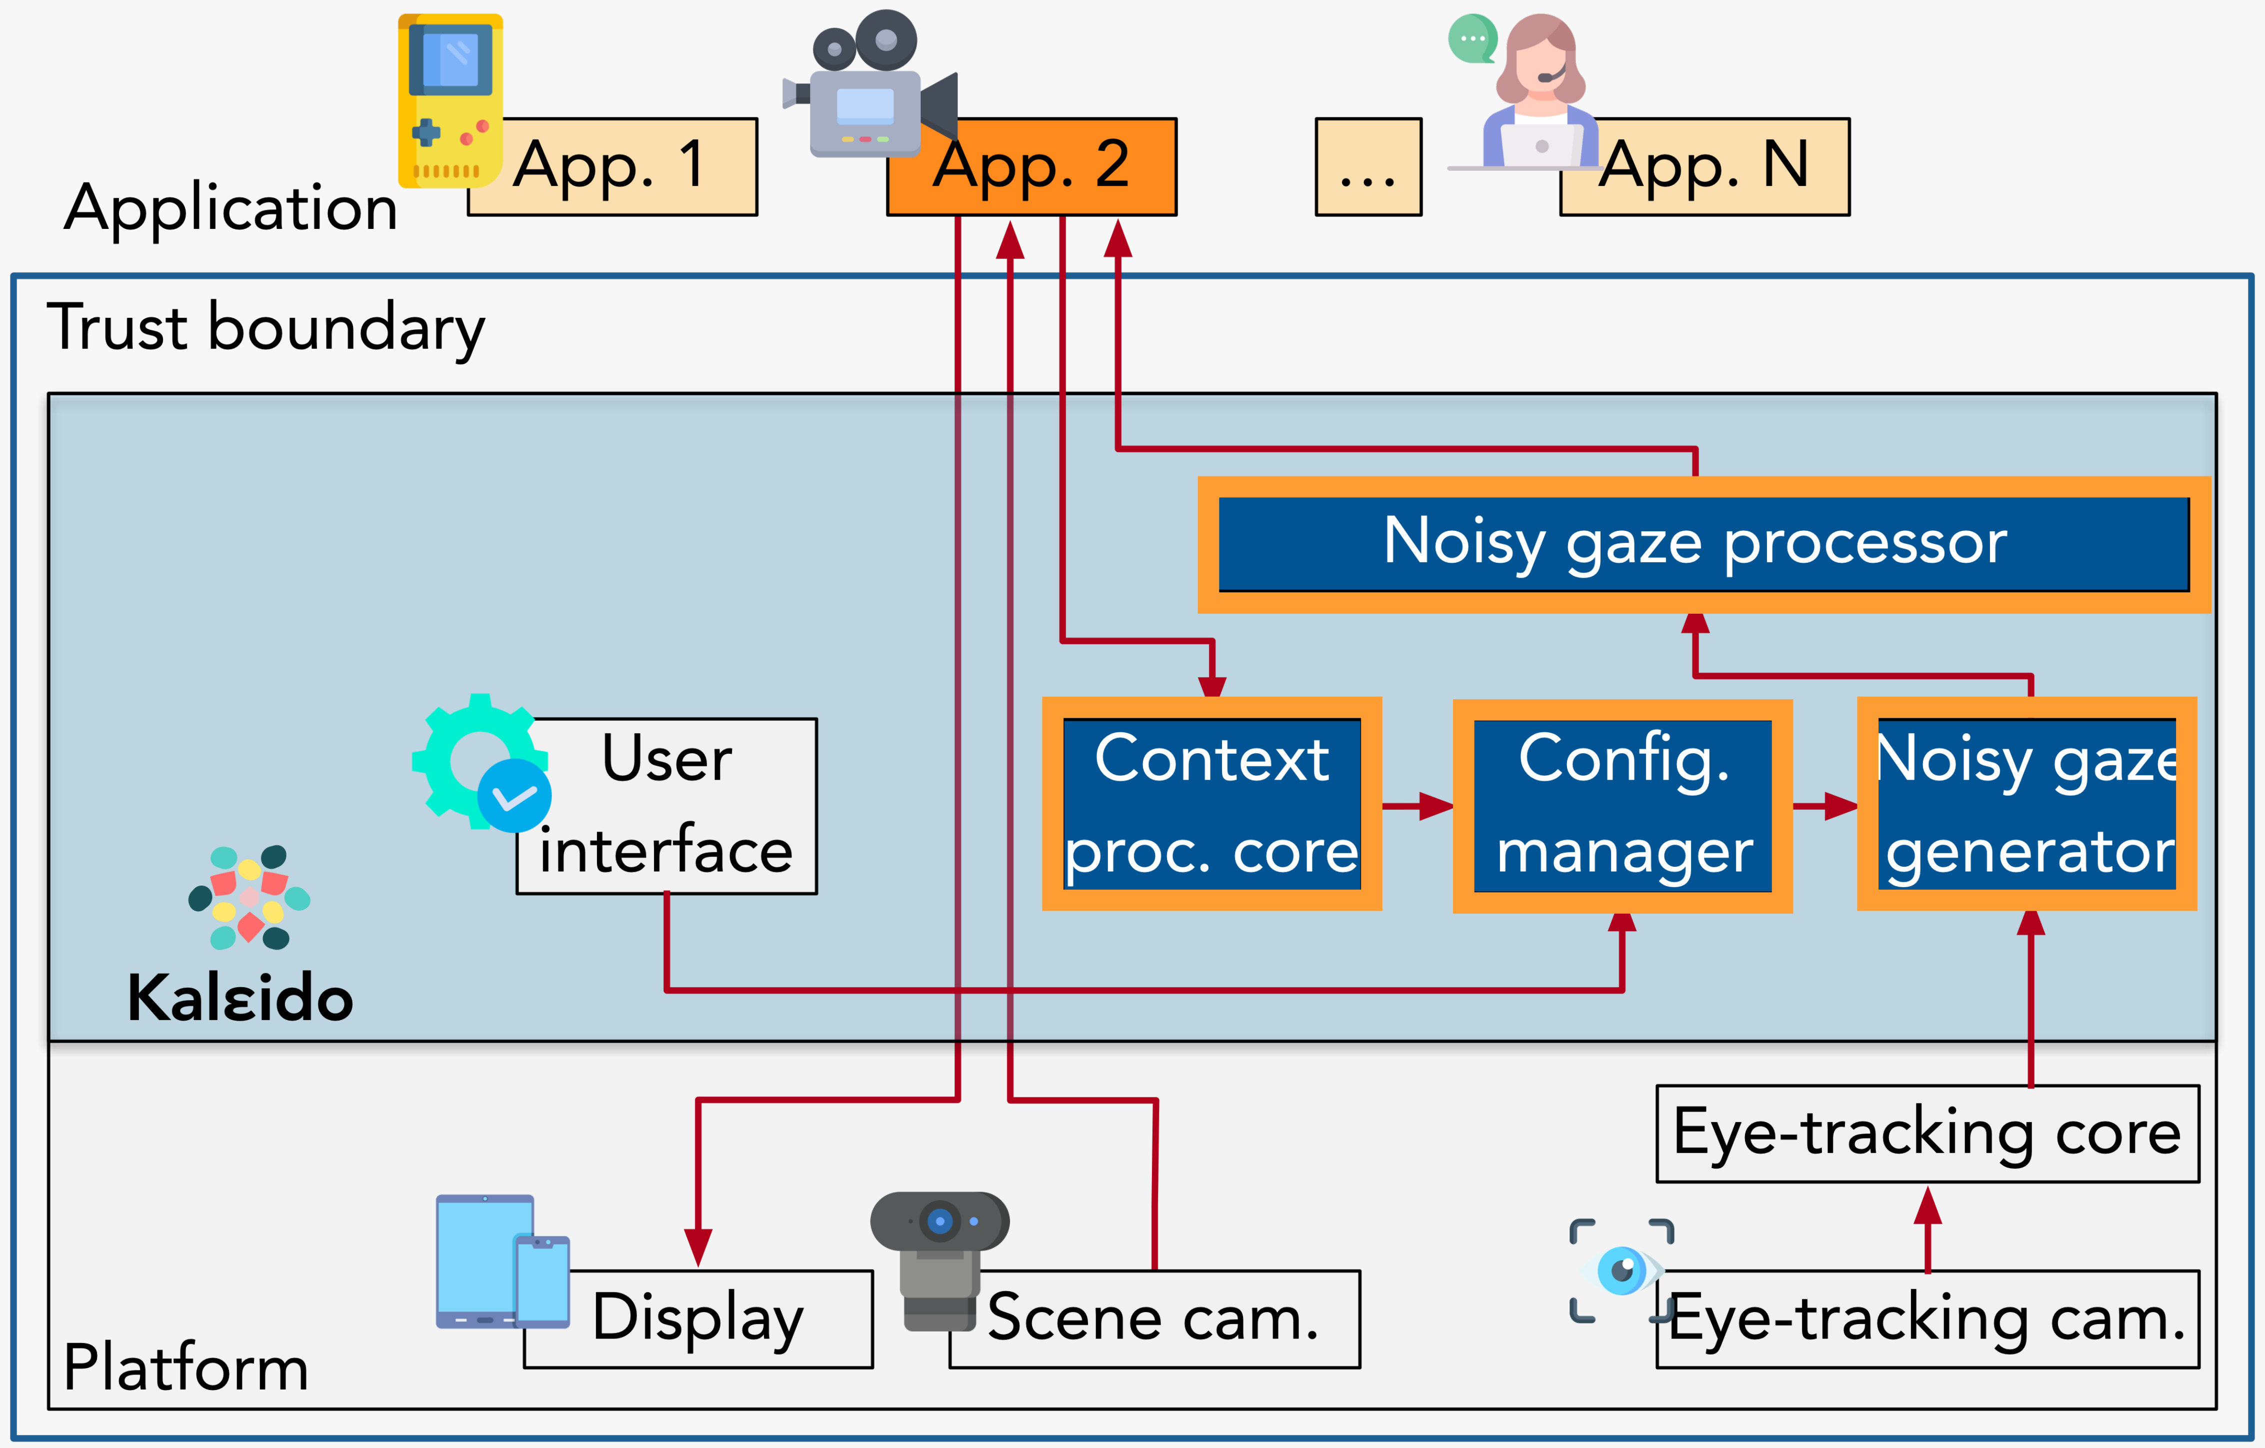

Config. manager configures privacy budget **ε**, window length *w*, and radius *r*

Context proc. core extracts ROI for setting *r*

Noisy gaze gen. noises each raw gaze online

Noisy gaze proc. allows local post-processing

KALεIDO: IMPLEMENTATION

No privacy (ε=∞)

Low privacy (ε=3)

High privacy (ε=0.5)

Evaluation Focus

User perception

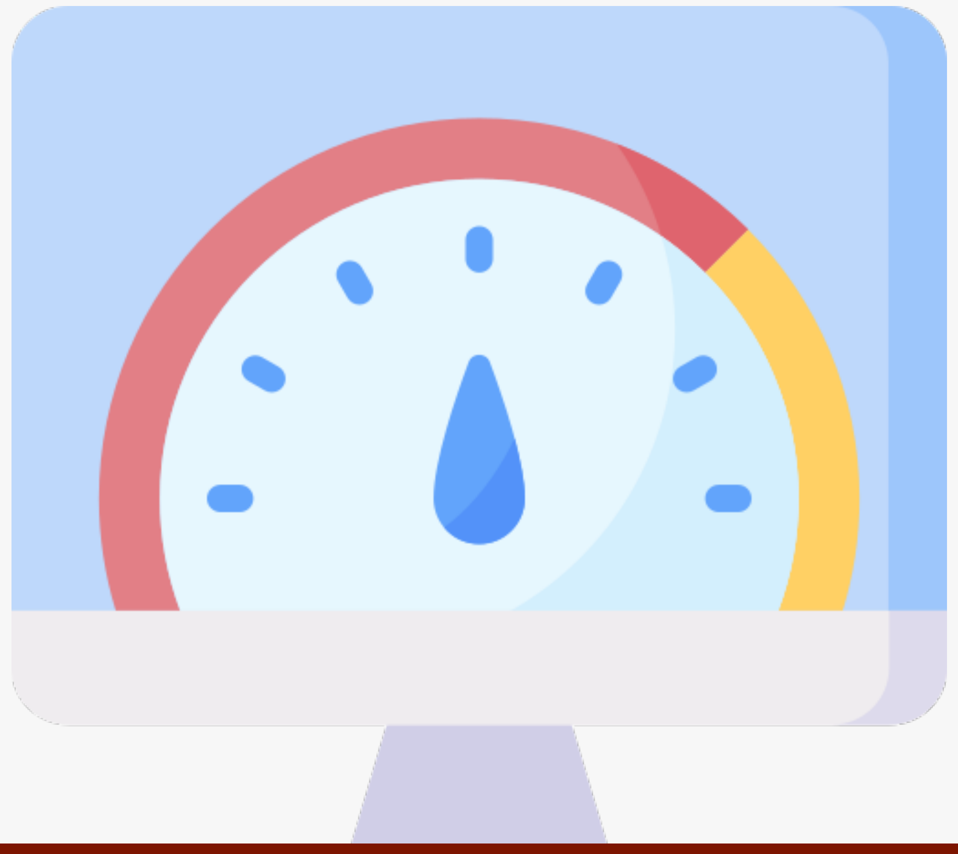

System performance

Effectiveness against attacks

User Study: Setup

- Remote user study with the PC webcam eye-tracking game (approved by our IRB)
- 11 users, each with a study session about 35 minutes in total
- Five settings evaluated in anonymized and randomized order except the control knob setup

User Study: Result

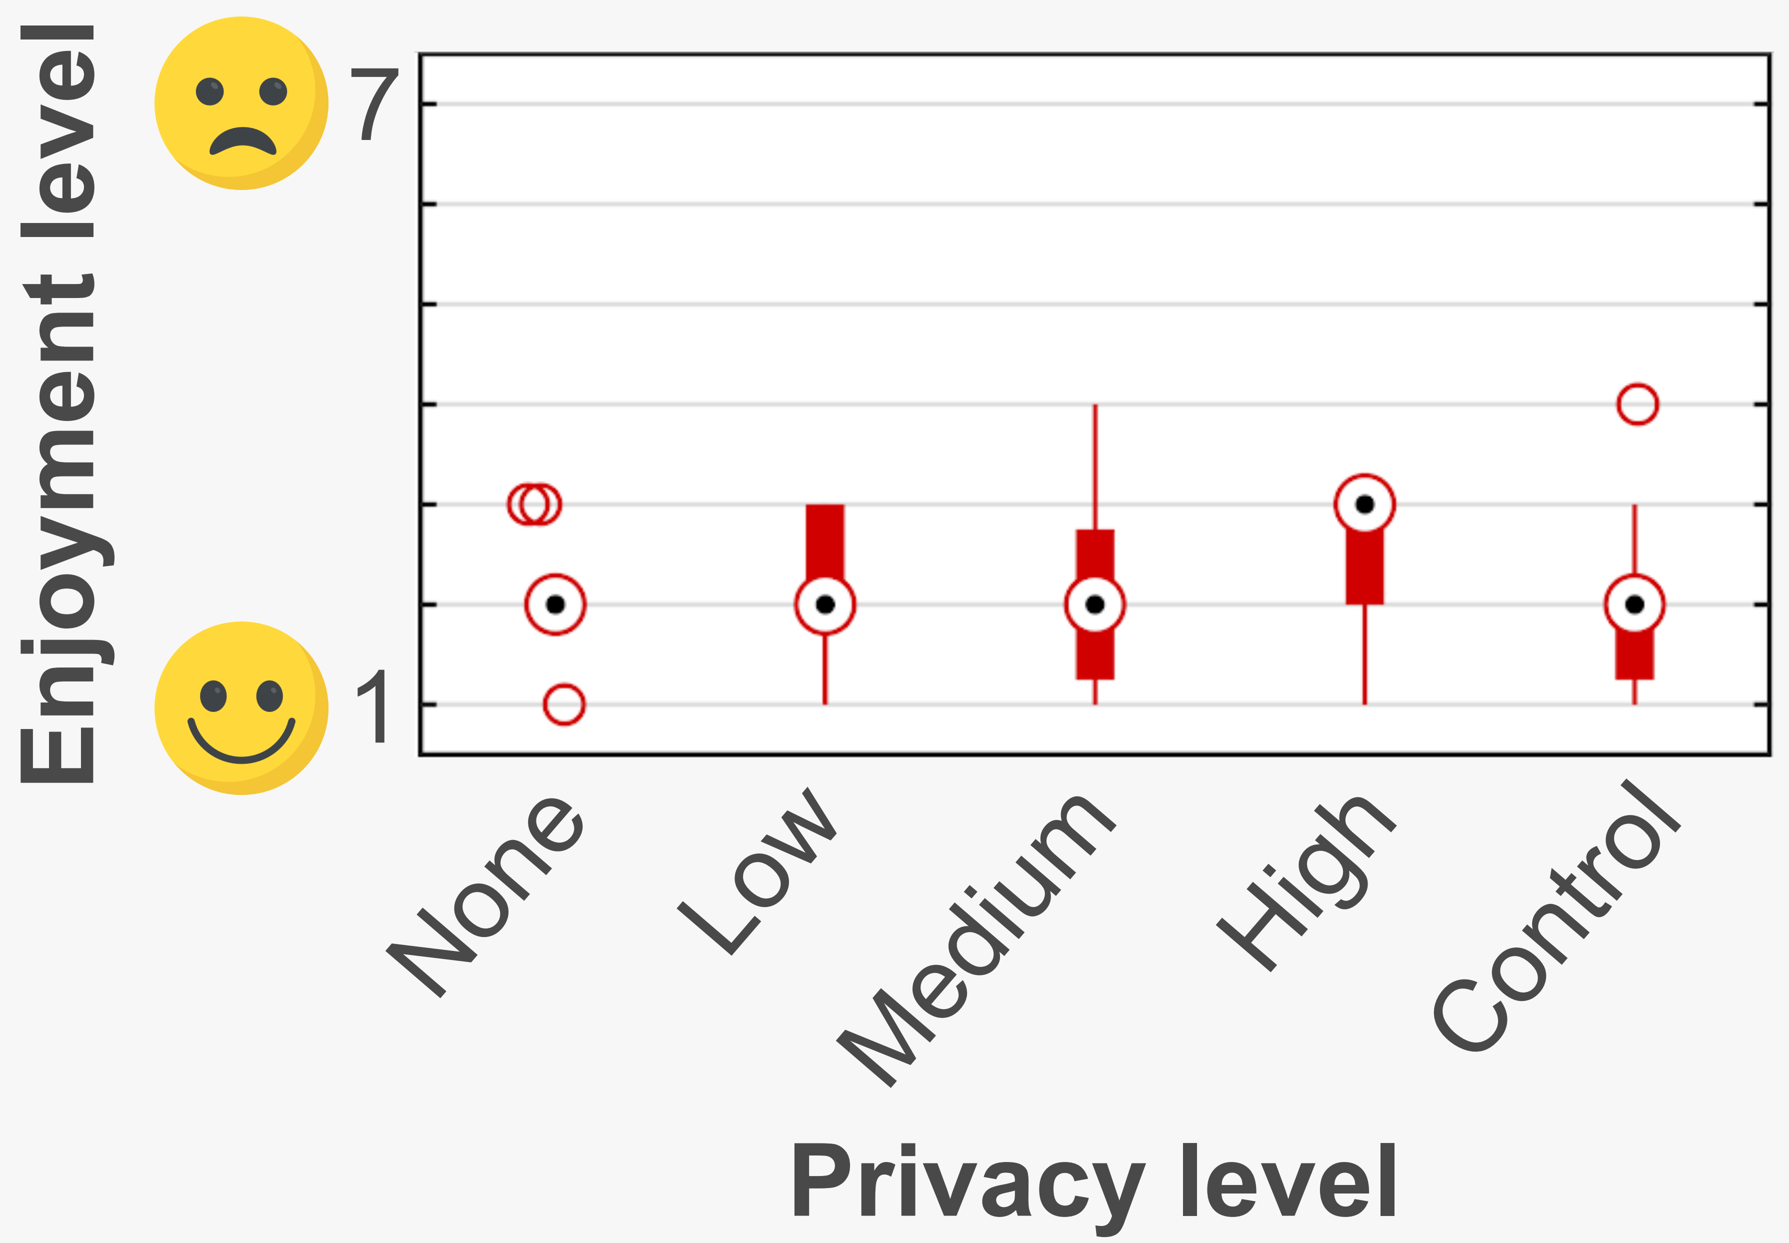

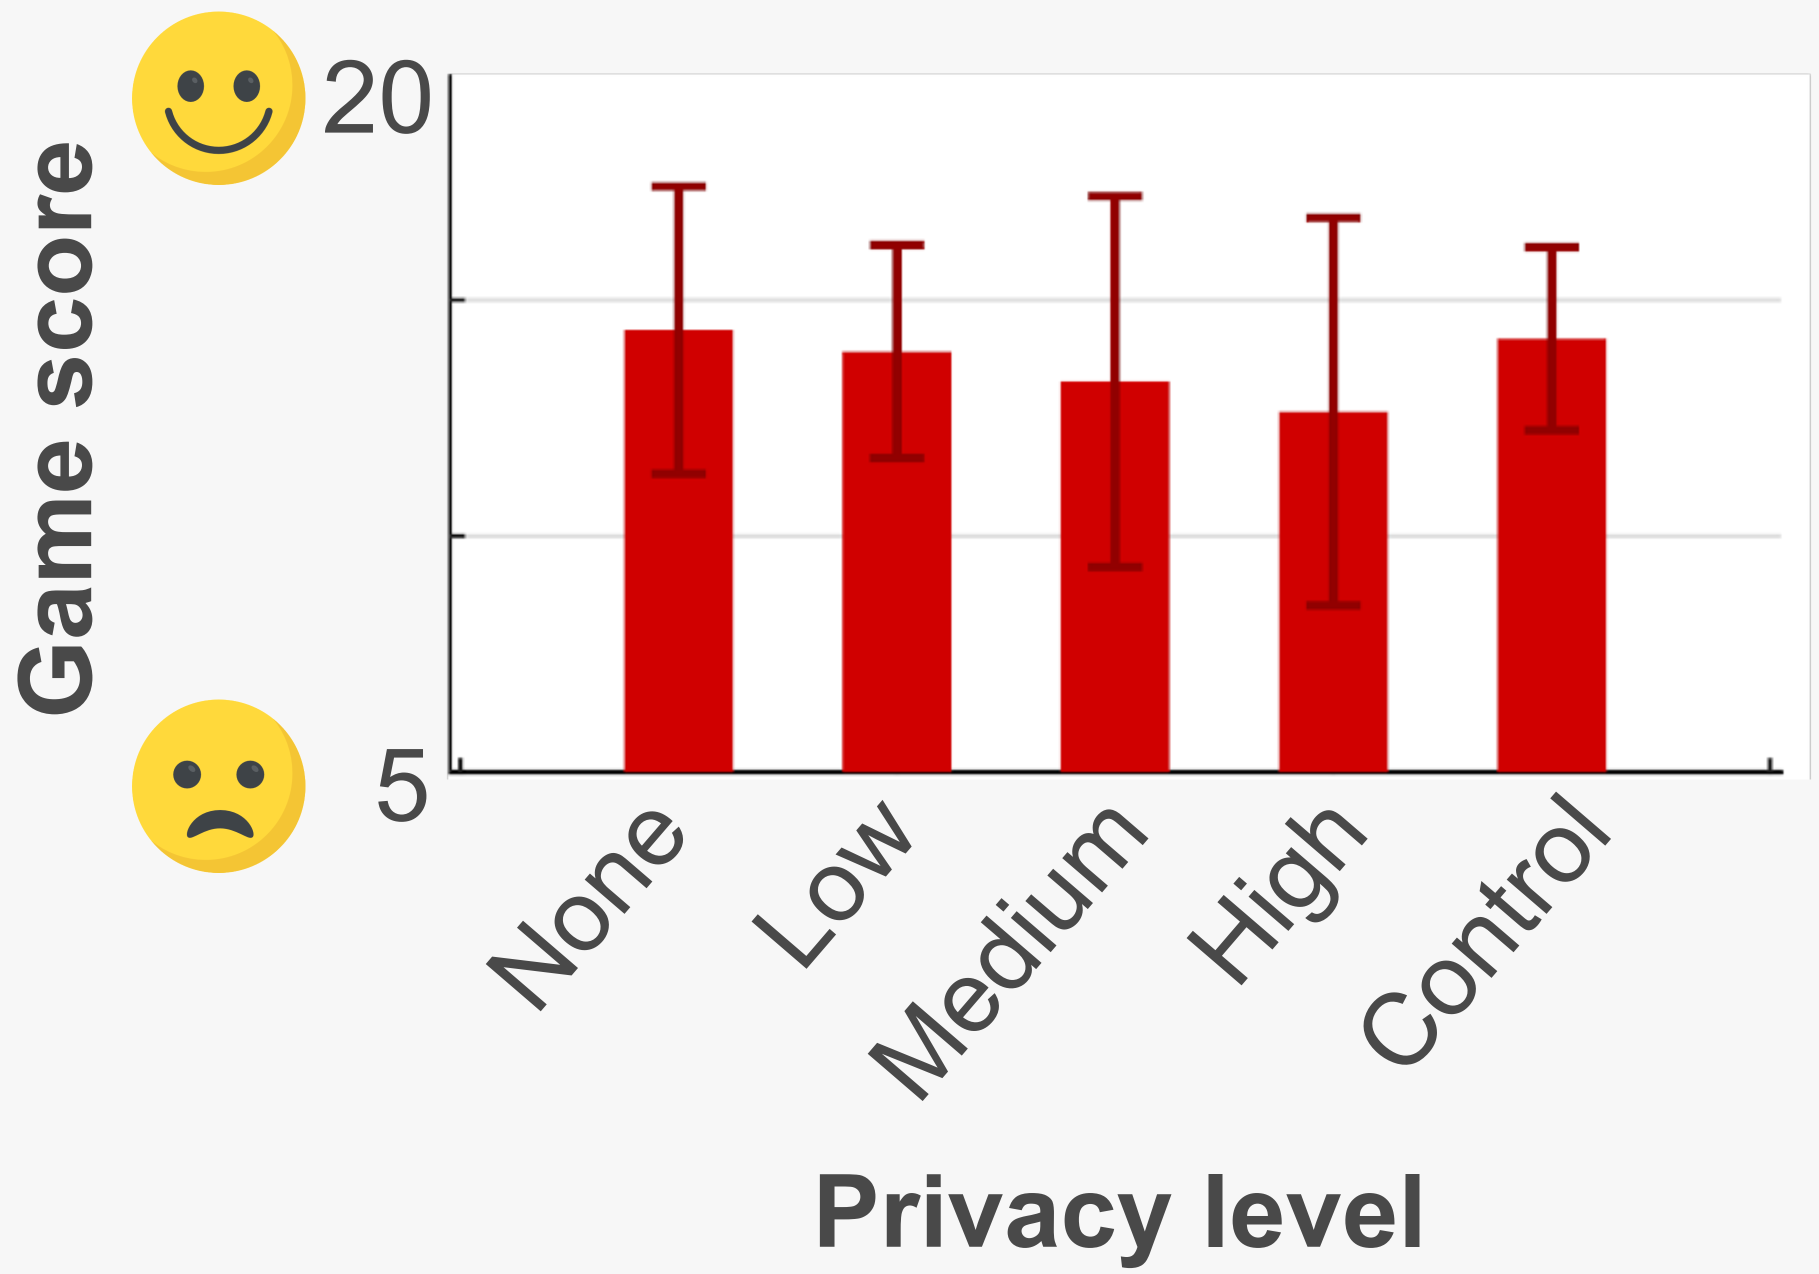

- **Metrics:** (1) subjective enjoyment level; (2) game score (# of rabbits taken)
- **Takeaway:** negligible experience degradation with low privacy; even high privacy poses minor impact

System Performance

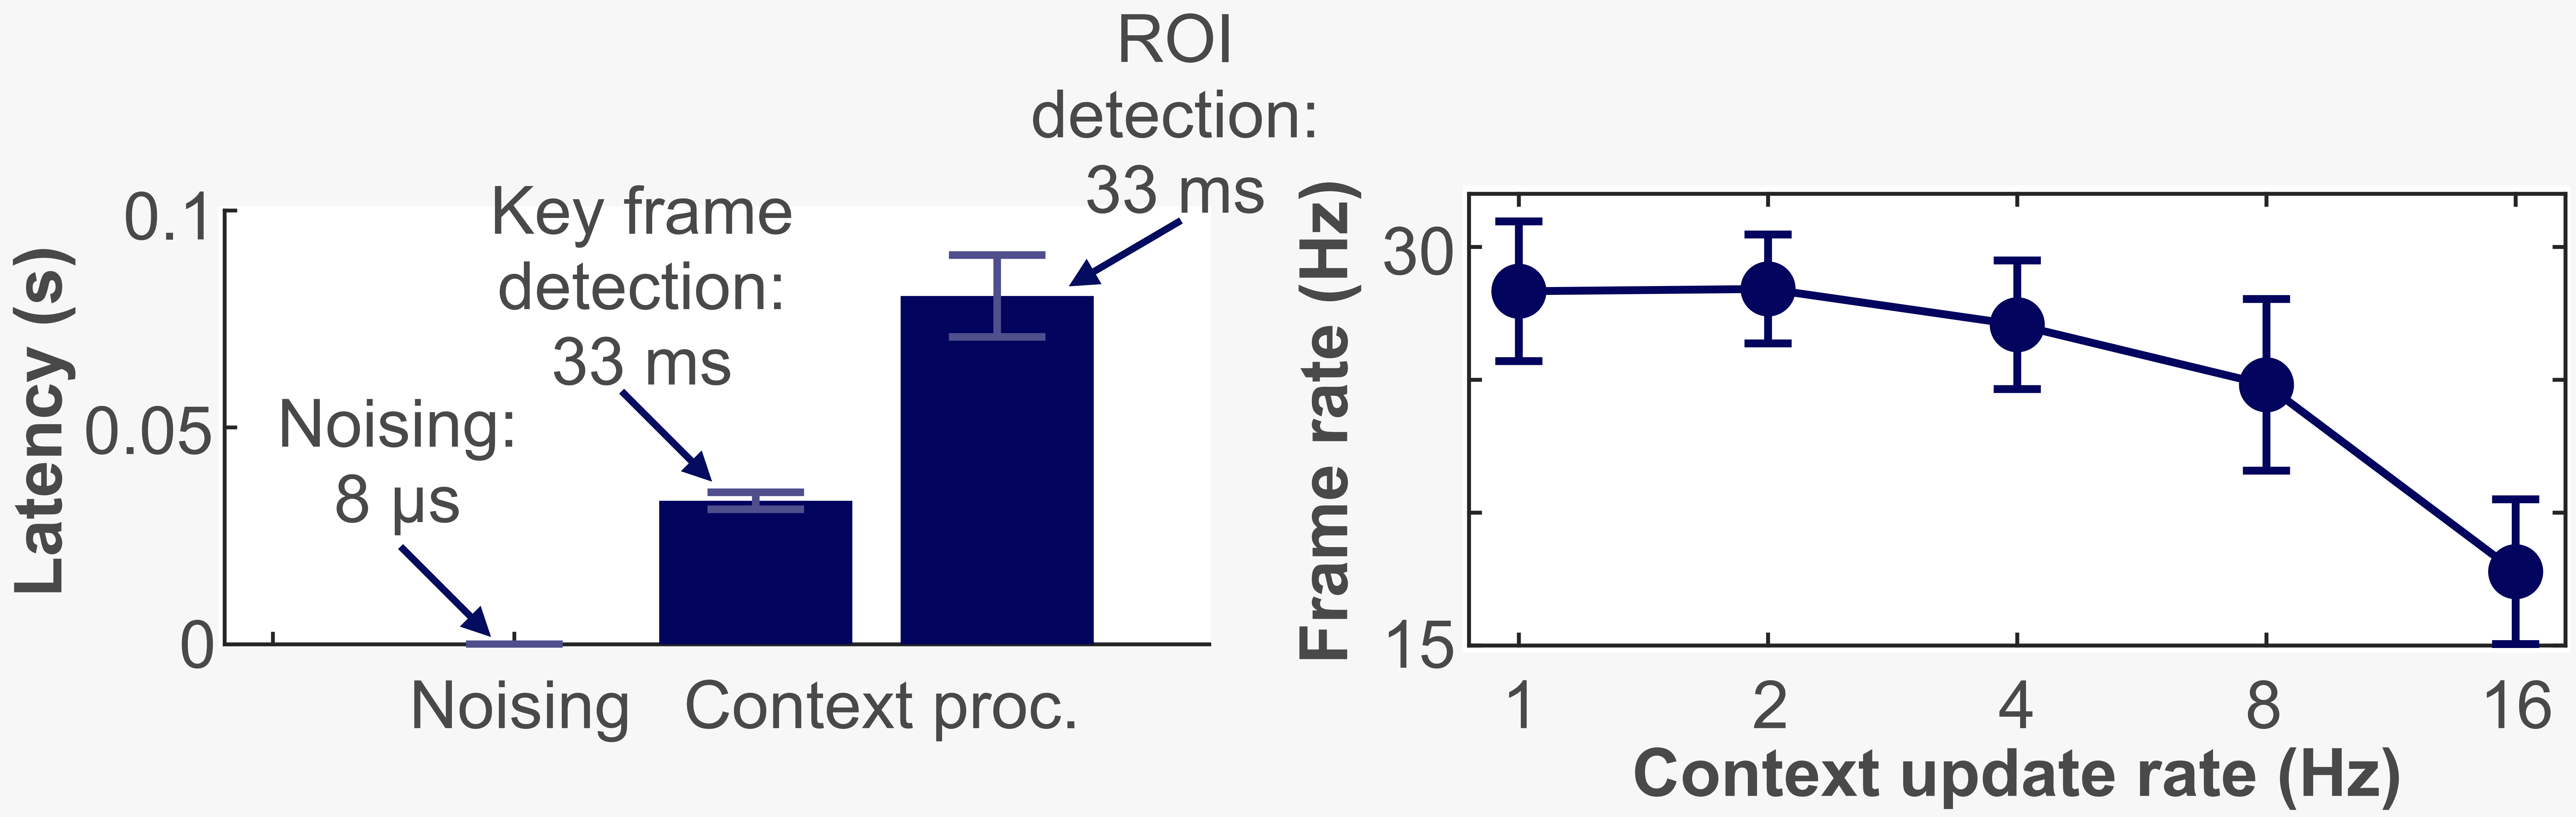

- **Platform:** Intel I7-7700 & Nvidia GTX1080
- **Takeaway:** noising takes negligible latency; performance not degraded greatly even at very frequent context processing rate of 8 Hz

Effectiveness Against Attack on Interest

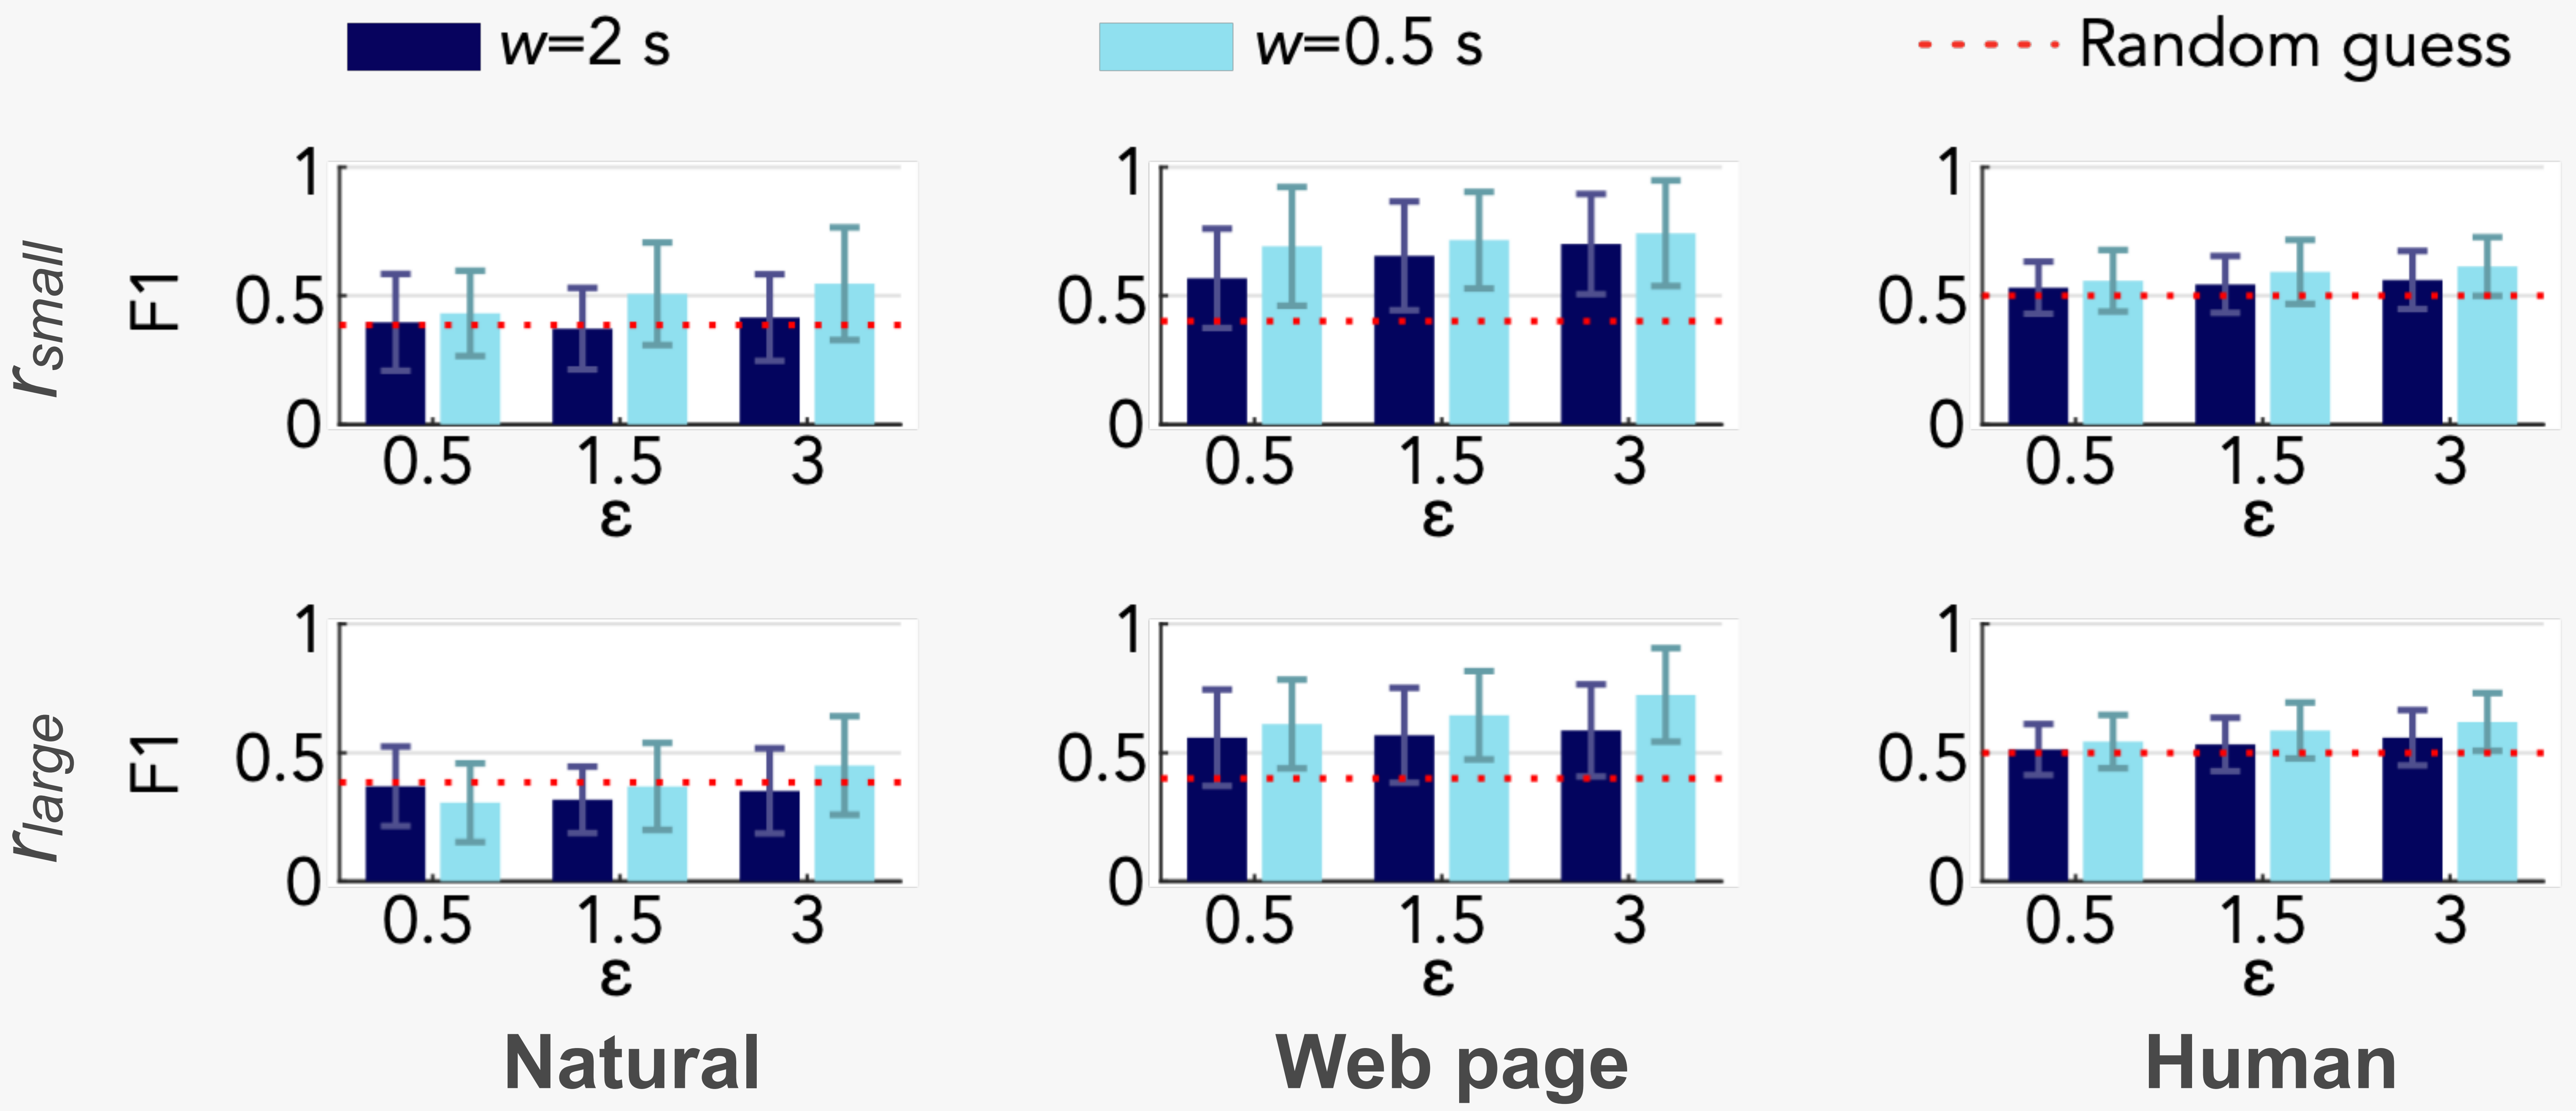

- **Dataset:** PC eye tracking for viewing 30 images (at least 19 users)
- **Attack setup:** identify users with distinct attention patterns per image by clustering
- **Takeaway:** attacker's success brought to random guess at high privacy; even lower privacy thwarts attacks greatly

Effectiveness Against Attack on Biometrics

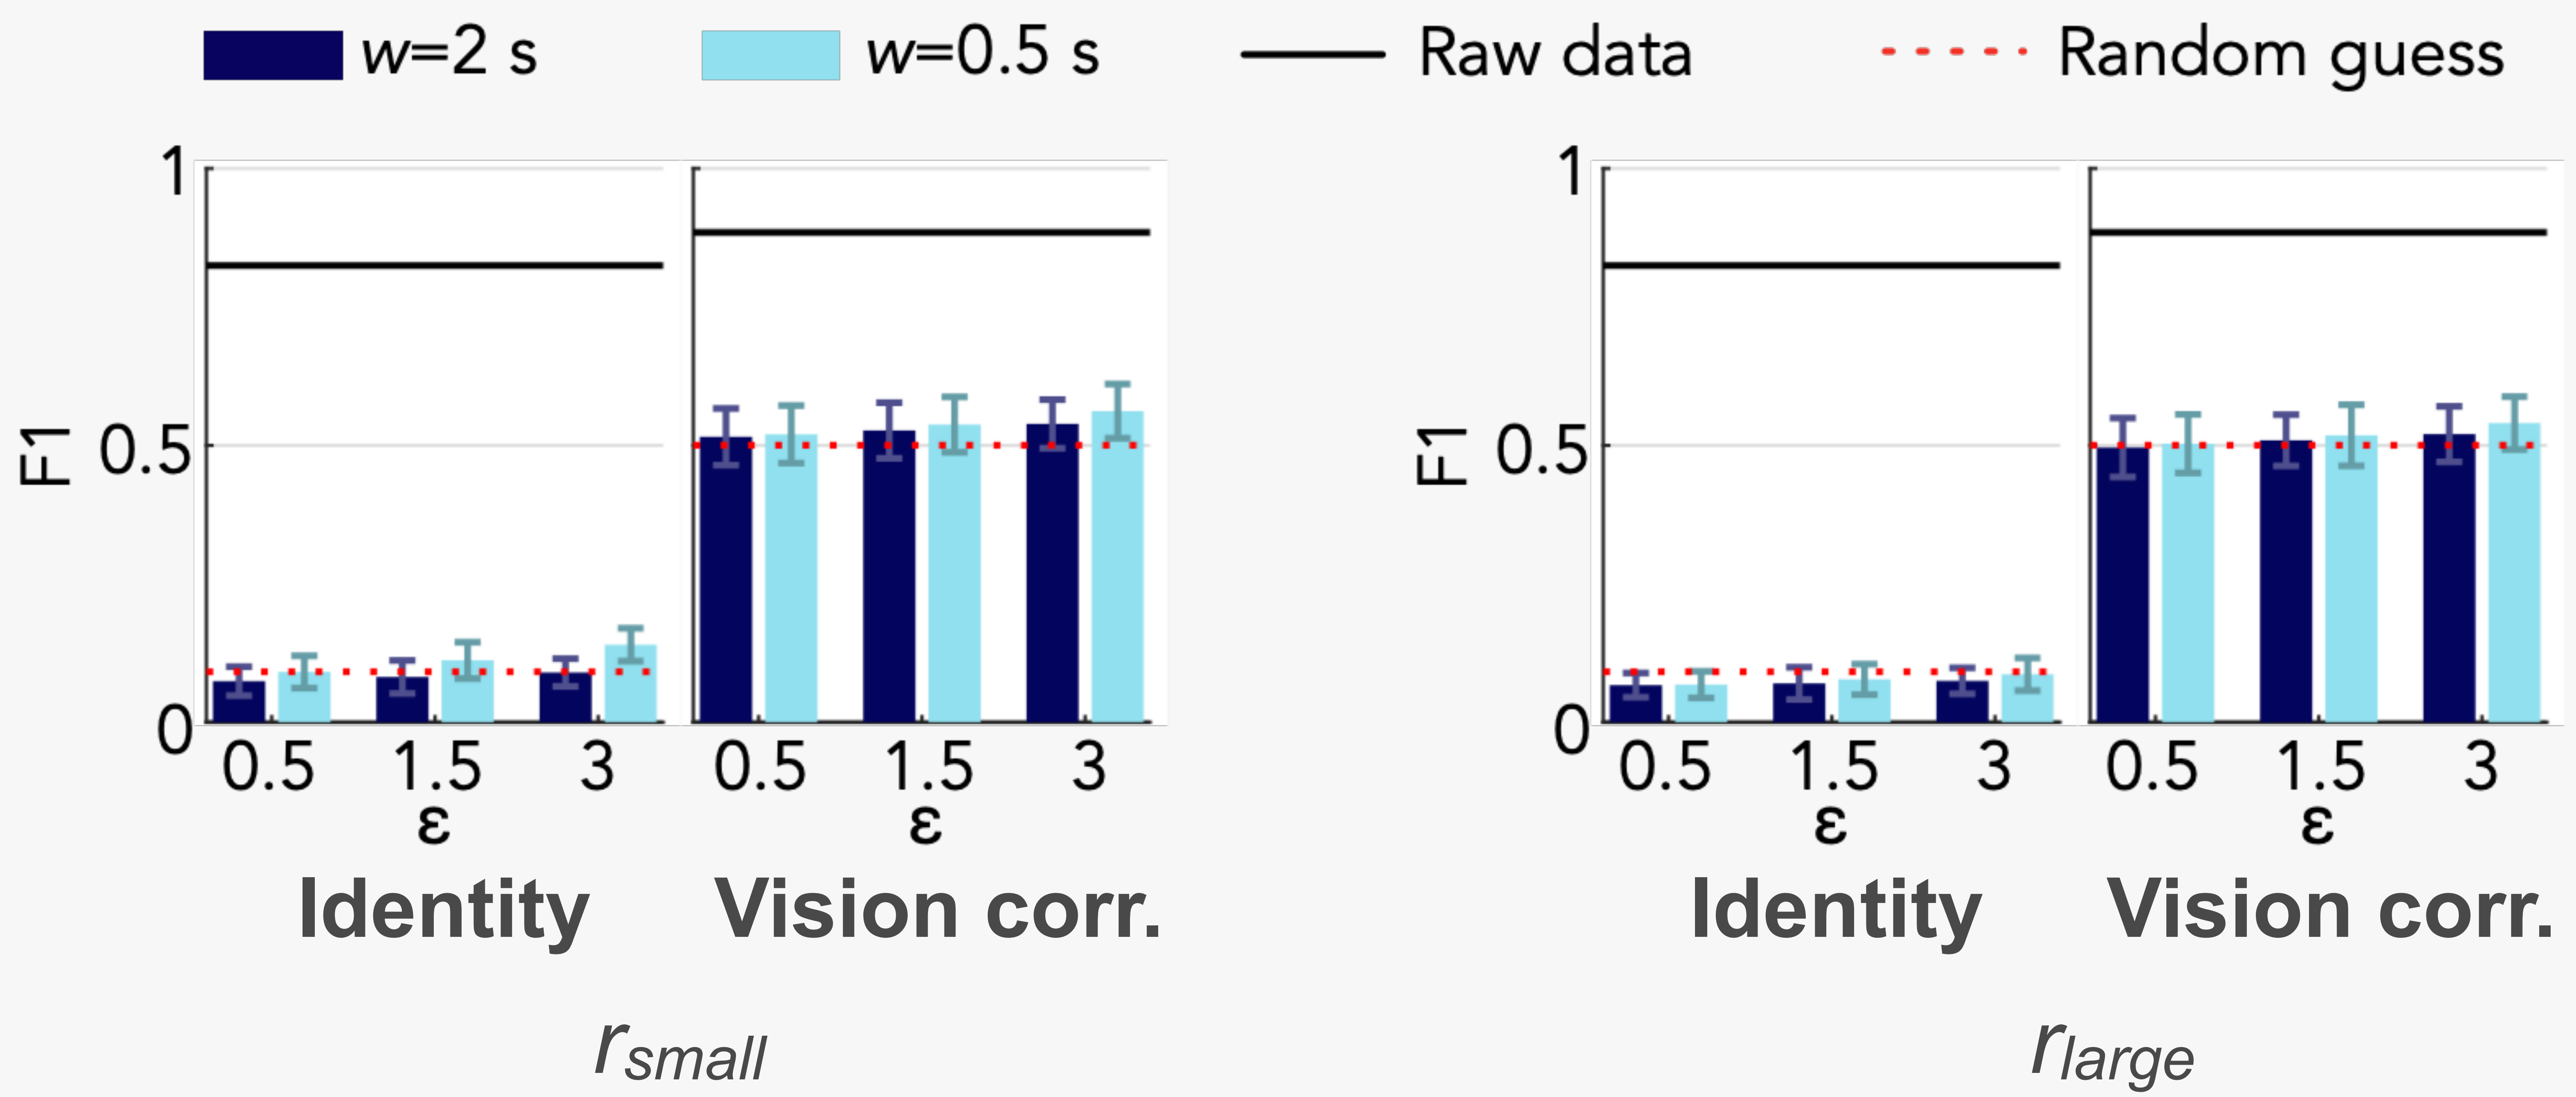

- **Dataset:** VR eye-tracking during video sessions for 12 unique videos with 11 users
- **Attack setup:** identify user traits by classifiers trained on biometric features
- **Takeaway:** attacker's success brought to random guess even with low privacy configuration for both traits

Conclusion

- Kalεido, the first system to protect privacy of real-time eye tracking
- Deploying differential privacy by leveraging semantics of eye gazes
- Seamlessly integration with existing eye-tracking ecosystems

Contact

Jingjie Li
Ph.D. candidate, UW-Madison
Research interest: human-centered computing, security and privacy
Email: jingjie.li@wisc.edu
Homepage: https://jingjieli95.github.io/

Reference

Hagestedt, I., Backes, M. and Bulling, A., 2020, June. Adversarial Attacks on Classifiers for Eye-based User Modelling. In Proceedings of the 12th ACM Symposium on Eye Tracking Research and Applications, pp. 1-3.

Steil, J., Hagestedt, I., Huang, M.X. and Bulling, A., 2019, June. Privacy-aware eye tracking using differential privacy. In Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications, pp. 1-9.

Andrés, M.E., Bordenabe, N.E., Chatzikokolakis, K. and Palamidessi, C., 2013, November. Geo-indistinguishability: Differential privacy for location-based systems. In Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security, pp. 901-914.

Kellaris, G., Papadopoulos, S., Xiao, X. and Papadias, D., 2014. Differentially private event sequences over infinite streams. In Proceedings of the VLDB Endowment, 7(12), pp.1155-1166.

Vancouver (photo by Aditya Chinchure), https://unsplash.com/photos/Bs_ac8eRkME

Basics, Tobii XR SDK, https://vr.tobii.com/sdk/learn/interaction-design/use-cases/social/basics/

Foveated rendering, Tobii, https://vr.tobii.com/foveated-rendering/

Nreal Light, https://venturebeat.com/2020/01/06/nreal-light-gets-nebula-3d-augmented-reality-ui-and-eye-tracking/