Privacy Preservation Strategies for Event-to-Image Reconstruction

A Comparative Study of Raw-Event Perturbation Strategies

Arda Tamgaç | Supervisor: Tunahan Parlayıcı | Responsible Professor: Nergis Tömen

Background & Problem

Event cameras are bio-inspired vision sensors that asynchronously detect per-pixel brightness changes, representing them as a sparse stream of events encoding time, location, and polarity [1, 2]. Unlike conventional cameras, they record no absolute intensity or colour, making them efficient and potentially *privacy-preserving*.

However, reconstruction models such as E2VID [3] can recover recognisable facial images directly from event streams, and the first event-based re-identification benchmark showed that identifying details survive reconstruction [4], undermining the privacy assumption. With event cameras being deployed in surveillance, autonomous vehicles, and human-computer interaction, their privacy implications have begun to receive research attention [5].

Research Question

What is the effect of applying **simple privacy-preserving perturbations** to event camera data in event-to-image reconstruction pipelines, in terms of **reconstruction quality** and **identifiability**?

Perturbation Methods

The following methods are implemented and evaluated across reconstruction quality and face identifiability, with a final comparison on the best privacy–utility tradeoff.

- **Polarity flipping** — randomly inverts the polarity of events with probability p
- **Spatial jitter** — displaces event coordinates by Gaussian noise $\mathcal{N}(0, \sigma^2)$
- **Event insertion & deletion** — each real event removed with probability ρ, an equal number of fake events inserted

Pipeline

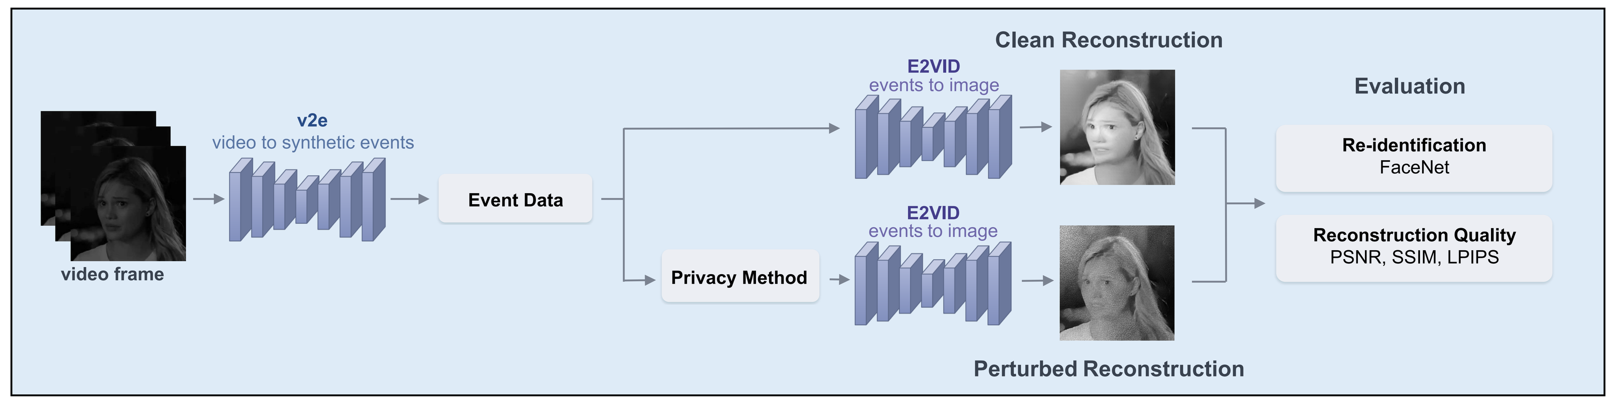

Figure 1. CelebV-HQ [6] clips are converted to events with v2e [7], perturbed at the raw-event level, reconstructed with E2VID [3], and evaluated for reconstruction quality and face identifiability.

Evaluation Metrics

- **Reconstruction Quality:** PSNR, SSIM, LPIPS
- **Identifiability:** FaceNet (Rank-1, identification rate, AUC)

Reconstructions & Face Re-identification

Representative reconstructions across perturbation methods and strength levels:

GT | Clean | p=0.3 | σ=3.0 | ρ=0.5

Figure 2. Ground truth, clean reconstruction, and the three perturbations.

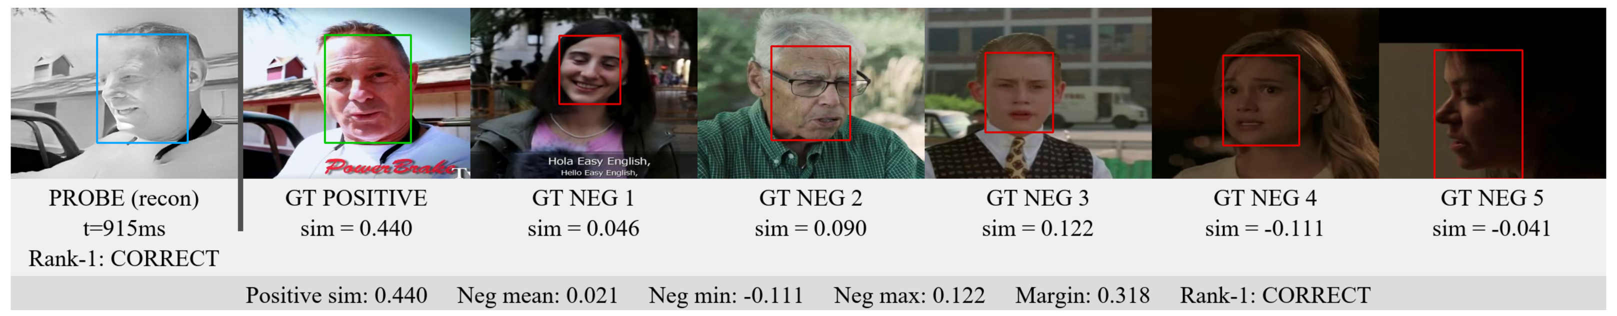

Figure 3. Baseline re-identification on a single clip. The reconstructed probe is correctly matched to the positive gallery entry over five negatives. Faces remain largely recoverable from unperturbed events.

Results

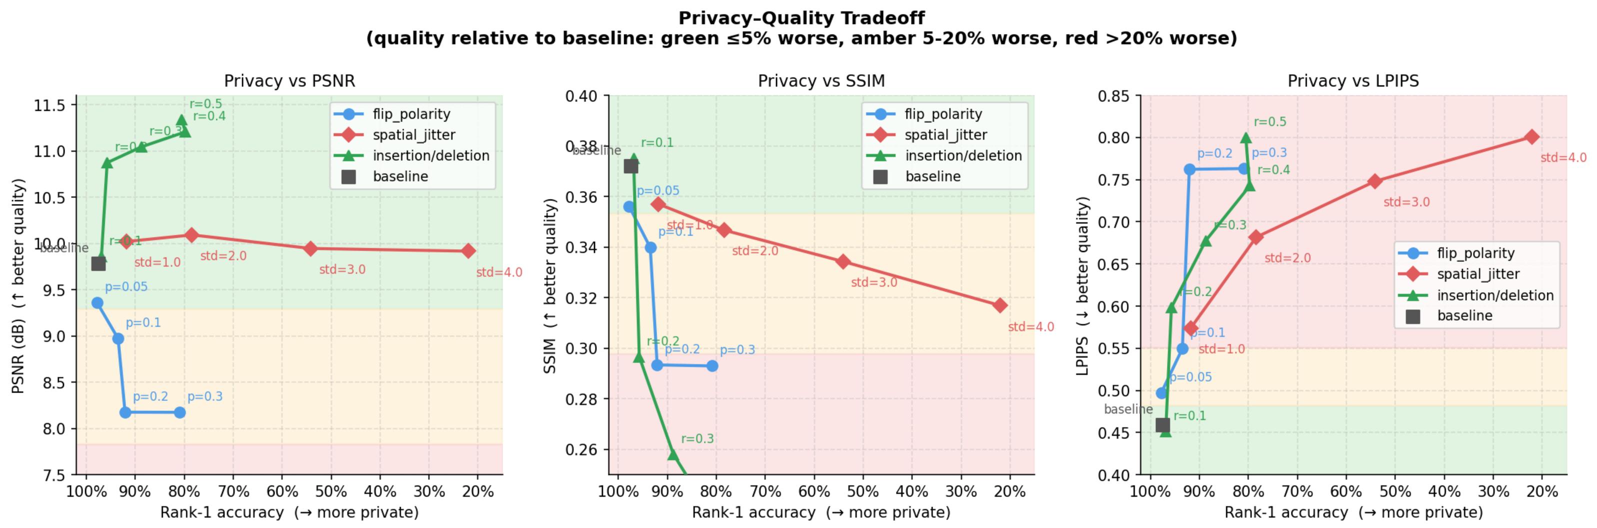

Figure 4. Privacy–utility tradeoff for all three quality metrics, plotted against Rank-1 accuracy. Spatial jitter (red) stays in or near the low-degradation bands while its Rank-1 accuracy falls furthest, making it the most favourable of the three; other methods leave the favourable bands at smaller privacy gains.

Spatial Jitter: Identifiability Drops While PSNR Stays Flat

Spatial jitter scrambles the spatial arrangement of facial features while leaving pixel-level quality largely intact: Rank-1 falls ~70 percentage points while PSNR stays essentially constant.

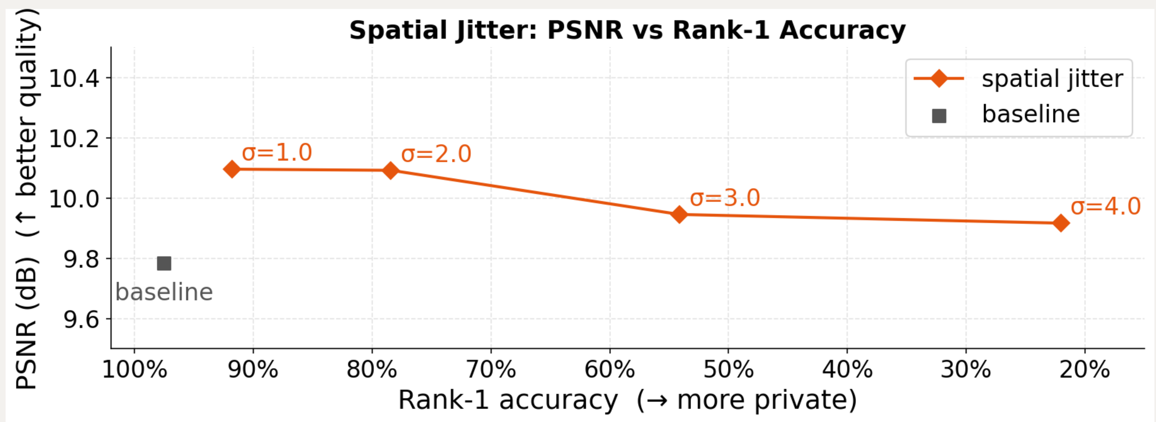

Figure 5. PSNR is flat across the jitter sweep even as Rank-1 accuracy collapses.

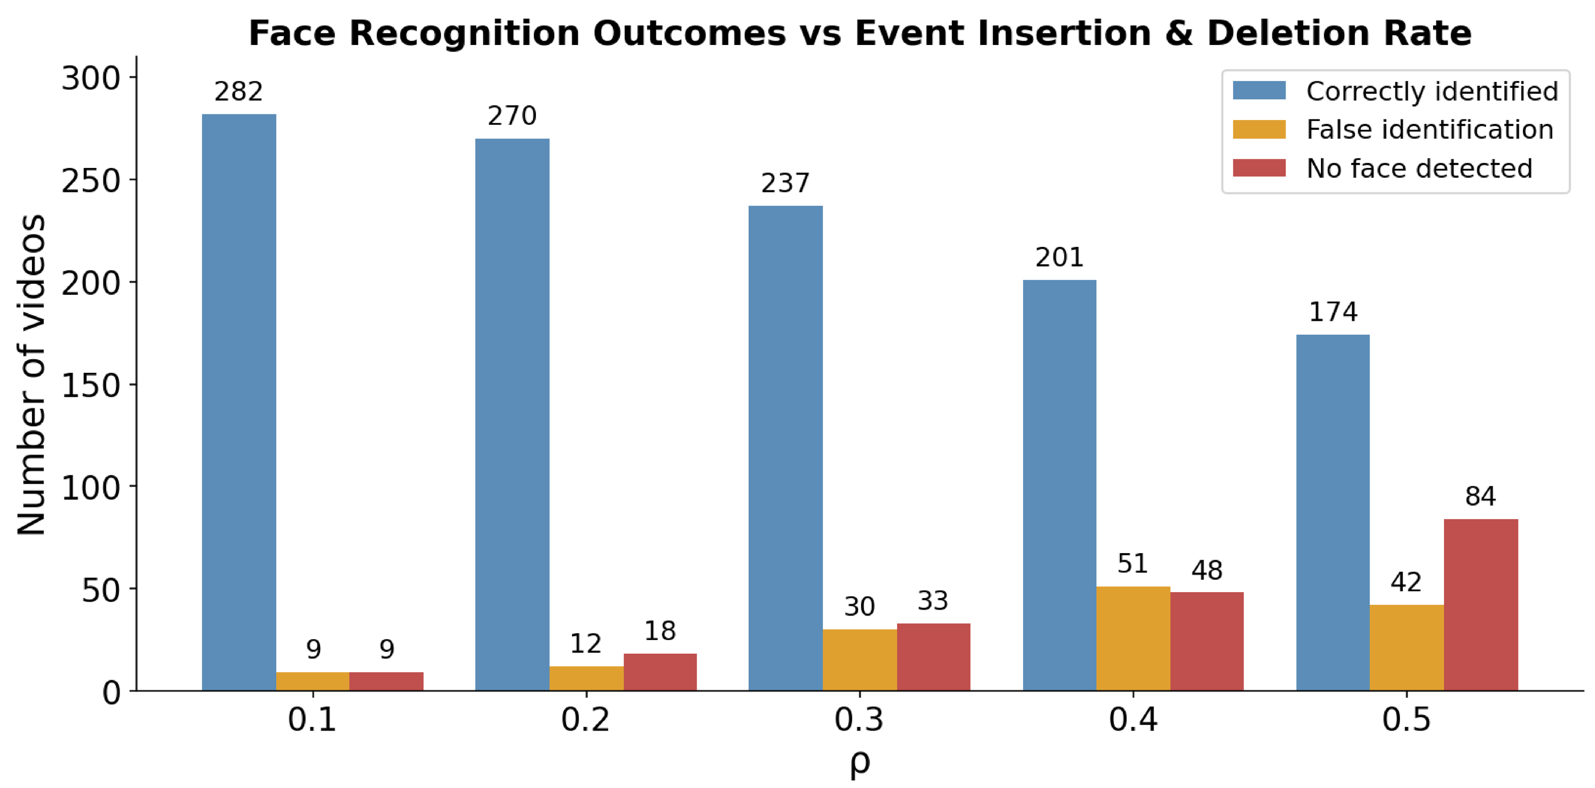

Figure 6. Rank-1 Face identification at different strength levels ρ

Method	Rank-1 (%)	Identified (%)	AUC	No-face (%)	PSNR (dB)↑	SSIM↑	LPIPS↓
Baseline	97.5	92.5	0.994	5.10	9.78	0.372	0.459
Polarity Flipping p=0.05	97.8	92.2	0.995	5.76	9.36	0.356	0.497
Polarity Flipping p=0.1	93.5	87.8	0.992	6.10	8.97	0.340	0.550
Polarity Flipping p=0.2	92.1	79.0	0.974	14.2	8.18	0.293	0.762
Polarity Flipping p=0.3	80.9	53.8	0.954	33.4	8.32	0.302	0.756

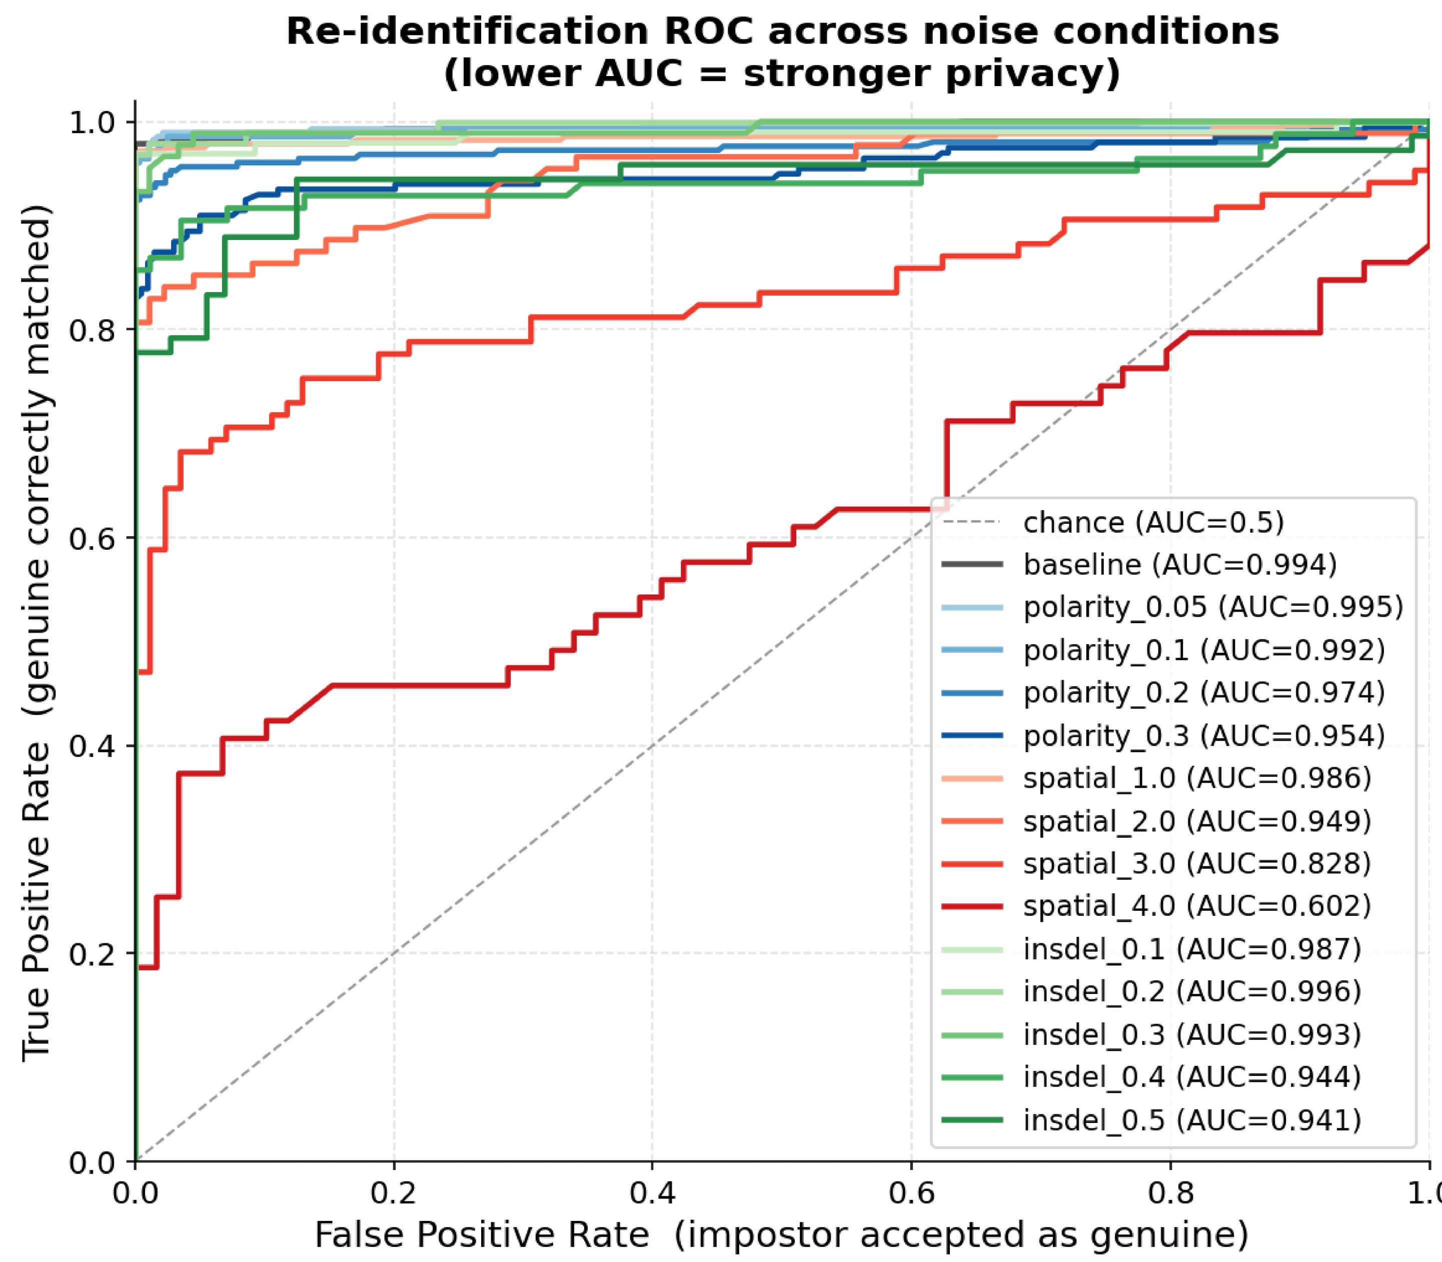

Figure 7. Re-identification ROC curves across all conditions. Almost all conditions stay near the top-left despite large drops in Rank-1 accuracy, showing identity remains separable under verification. Only the two strongest spatial-jitter settings bend toward chance.

Conclusions

- **Raw-event perturbations** were found to successfully reduce identifiability, but always at a direct cost to reconstruction quality.
- **Privacy gains** were found to be **real but shallow**, as identity remains statistically separable under verification (AUC), demonstrating that raw-event perturbations serve only as a **lightweight first layer** of defense rather than robust, complete anonymisation.
- **Spatial jitter at strength σ=3.0** was found to provide the optimal privacy–utility tradeoff, halving identifiability via *matcher confusion* while keeping reconstruction quality near baseline, whereas alternative strategies mostly destroy image readability completely.

References

[1] P. Lichtsteiner, C. Posch, and T. Delbruck. A 128×128 120 dB 15 μs latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits*, 43(2):566–576, 2008.

[2] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):154–180, 2022.

[3] H. Rebecq, R. Ranftl, V. Koltun, and D. Scaramuzza. High speed and high dynamic range video with an event camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(6):1964–1980, 2021.

[4] S. Ahmad, G. Scarpellini, P. Morerio, and A. Del Bue. Event-driven re-id: A new benchmark and method towards privacy-preserving person re-identification. In *WACV Workshops*, pages 459–468, 2022.

[5] K. Bendig, R. Schuster, N. Thiemer, K. Joisten, and D. Stricker. AnonyNoise: Anonymizing event data with smart noise to outsmart re-identification and preserve privacy. In *WACV*, pages 3159–3161, 2025.

[6] H. Zhu, W. Wu, W. Zhu, L. Jiang, S. Tang, L. Zhang, Z. Liu, and C. C. Loy. CelebV-HQ: A large-scale video facial attributes dataset. In *ECCV*, pages 650–667, 2022.

[7] Y. Hu, S.-C. Liu, and T. Delbruck. v2e: From video frames to realistic DVS events. In *CVPR Workshops*, pages 1312–1321, 2021.