Privacy Risks in Event Cameras — How Sensor Configuration Shapes Identity Inference

Author: Sofyan Ali (s.a.ali-2@student.tudelft.nl) • EEMCS, Delft University of Technology • **Supervisors:** Tunahan Parlayici, Prof. Nergis Tomen • **GitHub:** github.com/Stunner070/research-project-bsc-cse-tudelft

1 Motivation

- Event cameras record brightness changes instead of full frames, and are often seen as more privacy-friendly than traditional RGB video.
- However, recent literature indicates that event streams can still carry identity information (e.g., enabling event-based person Re-Identification).
- Sensor configurations dictate how events are triggered. We hypothesize that these configurations might strongly affect what information is recorded and how easily an identity can be inferred by an attacker.
- We utilize a frozen, state-of-the-art RGB face recognition model as standardized 'biometric ruler' to quantify how much recoverable identity information leaks through different hardware settings.

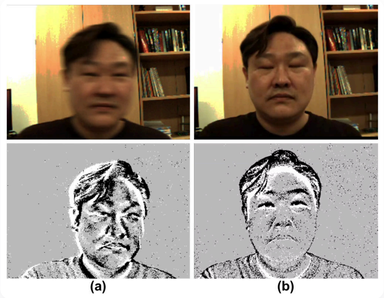

Figure 1: *Top row: RGB Camera | Bottom row: Event Camera*

2 Research Questions

> Main RQ: How does the ability of an attacker to infer identity from event-based representations vary across different event-camera sensor configurations, when the underlying scene content is fixed?

1. **Sub-RQ 1 (Hardware Analysis):** Which physically plausible event sensor parameters can be programmatically manipulated to regulate information density at the hardware layer?
2. **Sub-RQ 2 (Path A - Utility):** How do sensor parameter changes affect the spatial structure and visual quality of event-frame representations?
3. **Sub-RQ 3 (Path B - Privacy):** How well do sensor parameter changes reduce biometric identity leakage against a deep-learning attacker?

3 Methodology Overview

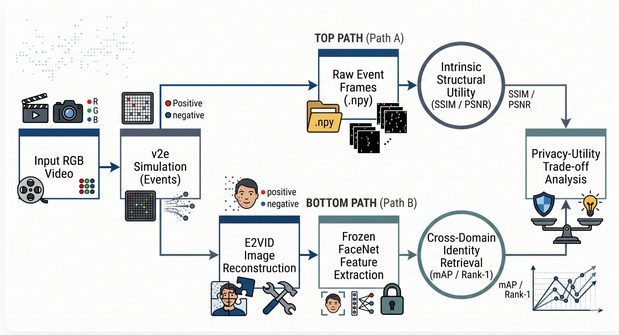

Figure 2: *Methodology Pipeline*

1. **Simulation:** Start with VoxCeleb2 videos; use v2e to generate Baseline vs. Adjusted event streams.
2. **Dual Path A (Intrinsic):** Compute SSIM and PSNR between Baseline and Adjusted .npy frames to measure structural utility loss.
3. **Dual Path B (Adversarial):** Reconstruct Adjusted events via E2VID and extract embeddings using a frozen FaceNet.
4. **Evaluation:** Perform Cross-Domain Retrieval—match 'Adjusted' queries against a 'Clean Baseline' gallery.

4 Experiment Setup

DATASET

- **VoxCeleb2:** Selected for its high identity consistency across multiple video tracks per subject. We use a canonical identity-disjoint split to ensure the attacker encounters unseen faces.

Figure 3: *VoxCeleb2 Dataset*

TABLE 1: SENSOR PARAMETERS (VARIED INDIVIDUALLY)

Parameter	Baseline	Test Setting
Spatial resolution	346 × 260	640 × 480
Contrast threshold	0.2, 0.2	0.1, 0.1
Photoreceptor bandwidth	200 Hz	100 Hz
Background leak event rate	0.1 Hz	5 Hz
Polarity asymmetry	0.2, 0.2	ON dominant (0.15, 0.25)

EVENT REPRESENTATION

Dual-Format Approach:

- **Raw .npy Event-Frames:** Chosen because they capture the immediate, integrated physical signal directly from sensor simulation without introducing network artifacts, serving as the pure baseline for structural utility.
- **E2VID Reconstructions:** Translates continuous asynchronous events into standard synchronous grayscale video. This represents the exact visual representation a real-world adversary generates to stage attacks.

ATTACKER MODEL

- **Frozen FaceNet:** Pre-trained on RGB photos, used as a worst-case 'black-box' adversary to measure biometric feature leakage.

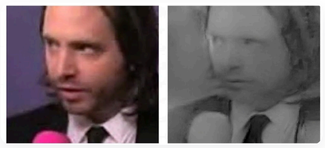

Figure 4:
Left: *Original frame from video*
Right: *Reconstructed grayscale image with E2VID*

5 Evaluation & Metrics

Intrinsic Utility Metrics (Video Quality):

- **SSIM (Structural Similarity):** Measures how well facial structures and edges are preserved.
- **PSNR (Peak Signal-to-Noise Ratio):** Measures the ratio of the true visual signal to the noise introduced by the sensor settings.

Adversarial Privacy Metrics (Attacker Risk):

- **Rank-1 Accuracy:** The probability that the attacker's top guess is the correct identity.
- **mAP (Mean Average Precision):** Evaluates the overall accuracy and quality of the attacker's ranked search results.

6 Results

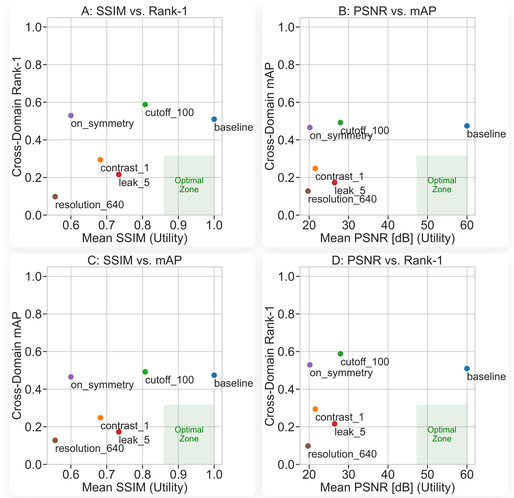

Figure 5: *Privacy-Utility Trade-off Space. The optimal privacy intervention should push the configuration into the high-utility, low-risk quadrant.*

- **Baseline Setup:** Yields a 50.98% Rank-1 attacker identification rate while establishing the maximum baseline for structural data utility.
- **Reduced Temporal Bandwidth (100 Hz):** Increases attacker success to 58.82% by filtering high-frequency noise, accompanied by only a minor decrease in structural utility.
- **Background Leak Injection (5 Hz):** Drops attacker success to 21.57% while successfully retaining moderate structural data utility.
- **Spatial Resolution Scaling (640x480):** Substantially decreases attacker identification accuracy while concurrently causing severe structural utility degradation due to a spatial mismatch.

7 Conclusions & Future Work

CONCLUSIONS

- **No inherent anonymity:** Raw event streams preserve enough spatial geometry for a standard facial recognition model to accurately identify subjects.
- **The privacy paradox:** Reducing temporal bandwidth acts as a hardware denoiser. This inadvertently helps attackers generate cleaner facial reconstructions and increases their success rate.
- **Noise is an effective filter:** Injecting uniform background leak events severely disrupts the attacker's reconstruction network, successfully masking identity while retaining moderate structural utility.
- **Motion dependency:** Static hardware settings cannot guarantee uniform privacy because effectiveness depends heavily on the subject's motion dynamics. A setting that protects a stationary subject may fail when they move.
- **Resolution scaling:** The massive drop in attacker success observed at higher resolutions is a spatial-mismatch artifact rather than a true physical privacy barrier.

FUTURE WORK

- **Physical Validation:** Test these findings on real hardware (e.g., DAVIS or DVXplorer sensors) to account for real-world optical, thermal, and manufacturing anomalies.
- **Advanced Threat Models:** Evaluate privacy leakage against custom, event-native ReID models instead of standard RGB networks to test true limits.
- **Adaptive Tuning:** Explore smart sensors that dynamically adjust noise rates or contrast thresholds in real-time based on scene activity.
- **Hybrid Systems:** Combine hardware-level filtering with software-based anonymization pipelines to securely balance privacy and operational utility.