Privacy Preservation in Event-Based Vision: Risks, Methods, and Trade-offs

The effect of applying perturbations on the privacy and visual naturalness of face images reconstructed from event-based data

Matei-Ioan Oprescu
m.i.oprescu@student.tudelft.nl
Supervisors: Nergis Tömen, Tunahan Parlayici
EEMCS, TU Delft CSE3000 Research Project

TU Delft

1 EVENT-BASED VISION

Event-based cameras capture brightness changes asynchronously as event streams.
Each pixel independently triggers an event only when its local brightness changes by a fixed threshold.
Because this raw event data is sparse, it was often considered privacy-preserving by design.

(a) Original Event-Camera

(b) E2VID Reconstruction

(c) E2VID Reconstruction, differenT thresholds

2 BACKGROUND AND MOTIVATION

Facial Recognition systems pose serious privacy and mass surveillance risks when deployed in the wild.

Privacy protection methods like AMT-GAN and Adv-CPG can be used to fool Face Recognition by adding makeup or altering face features.
These have been developed and applied to RGB images only.

Neural networks like E2VID can now reconstruct highly detailed, recognizable grayscale frames from event streams, recreating the privacy risk.

3 RESEARCH QUESTION

How do event-based reconstruction parameters and the choice of generative adversarial protection method determine the privacy-naturalness trade-off on event-reconstructed face images?

RQ1: Does AMT-GAN's adversarial signal transfer from clean RGB to event-reconstructed inputs?
RQ2: How do event-simulation parameters specifically the v2e contrast threshold) affect protection effectiveness?
RQ3: Where does any observed failure originate? In the event reconstruction, the protection method, or the FR evaluator?
RQ4: How does the domain shift of event reconstruction impact the visual naturalness and structural diversity of diffusion-based protections like Adv-CPG?

4 PIPELINE AND METHODOLOGY

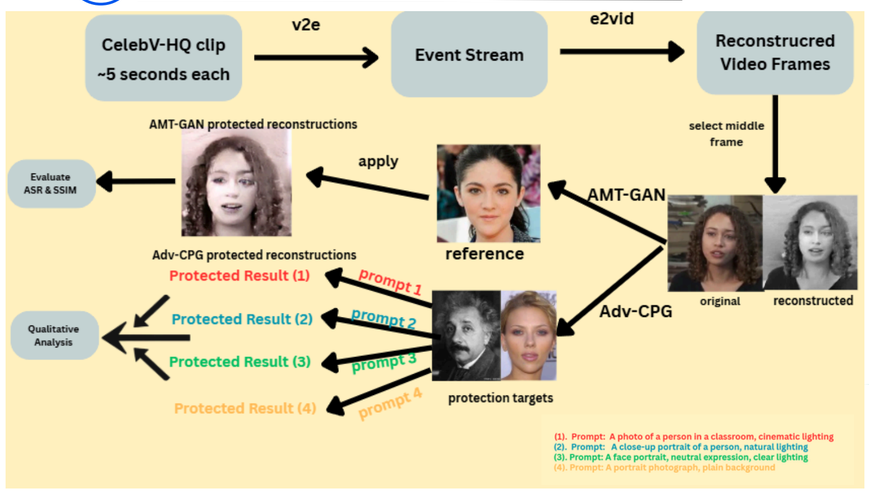

Dataset: 200 CelebV-HQ face videos.
Simulation: Videos converted to event streams using V2E simulator.
Reconstruction: Event streams rebuilt into grayscale face images using E2VID under different thresholds.
Protection: Applied AMT-GAN (nine makeup reference styles) and Adv-CPG (two target identities and four scene prompts).
Evaluation: Measured via Attack Success Rate (ASR) against four white-box FR models (ir152, irse50, facenet, mobile_face) and Structural Similarity Index (SSIM).

5 RESULTS

AMT-GAN

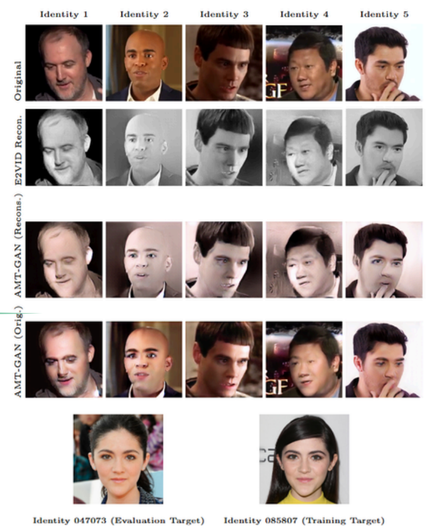

AMT-GAN baseline on original, unreconstructed, unprotected images.

FR Model	ASR @ FAR=0.1	ASR @ FAR=0.01	ASR @ FAR=0.001
ir152	36.38%	12.79%	4.49%
irse50	34.09%	10.81%	3.58%
facenet	40.11%	12.56%	0.61%
mobile_face	41.55%	15.45%	2.97%
Mean	**38.03%**	**12.90%**	**2.91%**

AMT-GAN protection results on event- reconstructed CelebV-HQ faces, evaluated using the official AMT-GAN protocol.under different reconstruction thresholds.

Threshold	n	Mean SSIM	ir152	irse50	facenet	mobile_face	Mean ASR
0.10	135	0.499	8.89%	5.93%	5.93%	5.93%	**6.67%**
0.15	138	0.547	7.97%	5.07%	6.52%	5.80%	**6.34%**
0.20	139	0.573	7.19%	3.60%	4.32%	6.47%	**5.40%**
0.25	130	0.578	6.15%	3.85%	7.69%	3.85%	**5.39%**
0.30	131	0.592	6.11%	4.58%	4.58%	3.05%	**4.58%**

AMT-GAN drops mean ASR to 5.64% when applied to reconstructed images. This is a 56% drop compared to the 12.90% original RGB baseline ASR.

Contrast threshold ablation (varying from 0.10 to 0.30) showed ASR remains consistently low (between 4.58% and 6.67%)

SSIM artificially rises at higher thresholds because smoother reconstructions contain less structural detail.

Protection reduction is caused by the event-to-pixel conversion.

Adv-CPG

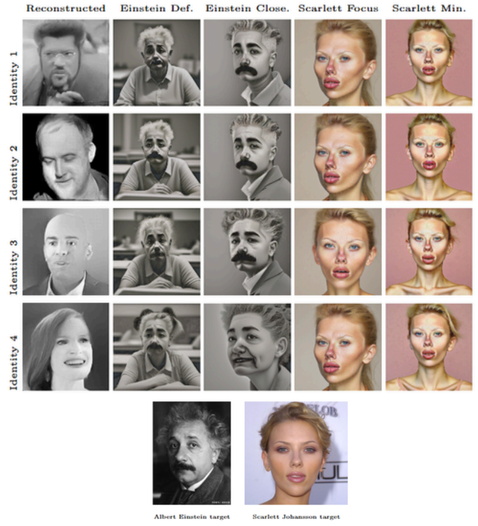

Qualitative study tested on a 10-image subset targeting Albert Einstein and Scarlett Johansson.

Adv-CPG generated unnatural artifacts and exaggerated facial features.

The model suffered a severe identity over-shift, making the original source visually unrecognizable.

Target prompts triggered a serious mode collapse, completely erasing the structural diversity of the source image.

Advanced privacy models fail because they expect sharp, high resolution photos, not the low-detailed reconstructions.

False Acceptance Rate (FAR): measures the likelihood of a identity verification system mistakenly granting access to an unauthorized user.
Attack Surface Rate (ASR): Measures the percentage of images that successfully fool the facial recognition system.
Structural Similarity Index (SSIM): Measures how visually similar and natural the protected image looks compared to the original

6 CONCLUSION & FUTURE WORK

Event-camera bottlenecks do not completely mitigate the perturbations, but significantly reduce them.

Current RGB-based adversarial protections are highly sensitive to domain shifts.

Applying existing RGB protections to event-reconstructed faces fails to provide appropriate privacy.

Changing event density parameters (reconstruction threshold) doesn't significantly change the protection's efficiency.

Reliable privacy protection in event-based vision requires designing adversarial protections built specifically for the event domain.