TeD-SPAD: Temporal Distinctiveness for Self-supervised Privacy-preservation for video Anomaly Detection

Joseph Fioresi, Ishan Rajendrakumar Dave, Mubarak Shah
University of Central Florida

https://joefioresi718.github.io/TeD-SPAD_webpage/

UCF CENTER FOR RESEARCH IN COMPUTER VISION

ICCV23 PARIS

Contributions

- ❖ Introduce privacy-preservation to VAD, show privacy leakage issue in existing works.
- ❖ Propose TeD-SPAD, a self-supervised privacy-preservation framework.
- ❖ Use of a **temporally-distinct triplet loss** to make the anonymization process more suitable for VAD.
- ❖ TeD-SPAD outperforms prior methods across all privacy-VAD benchmarks. On UCF-Crime, it eliminates **32.25%** privacy leakage with a 3.96% reduction in AUC performance.

Private Attribute Preservation

References

[1] Ishan Rajendrakumar Dave, Chen Chen, and Mubarak Shah. Spact: Self-supervised privacy preservation for action recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022.

Code:

Framework and Temporally-distinct Triplet Loss

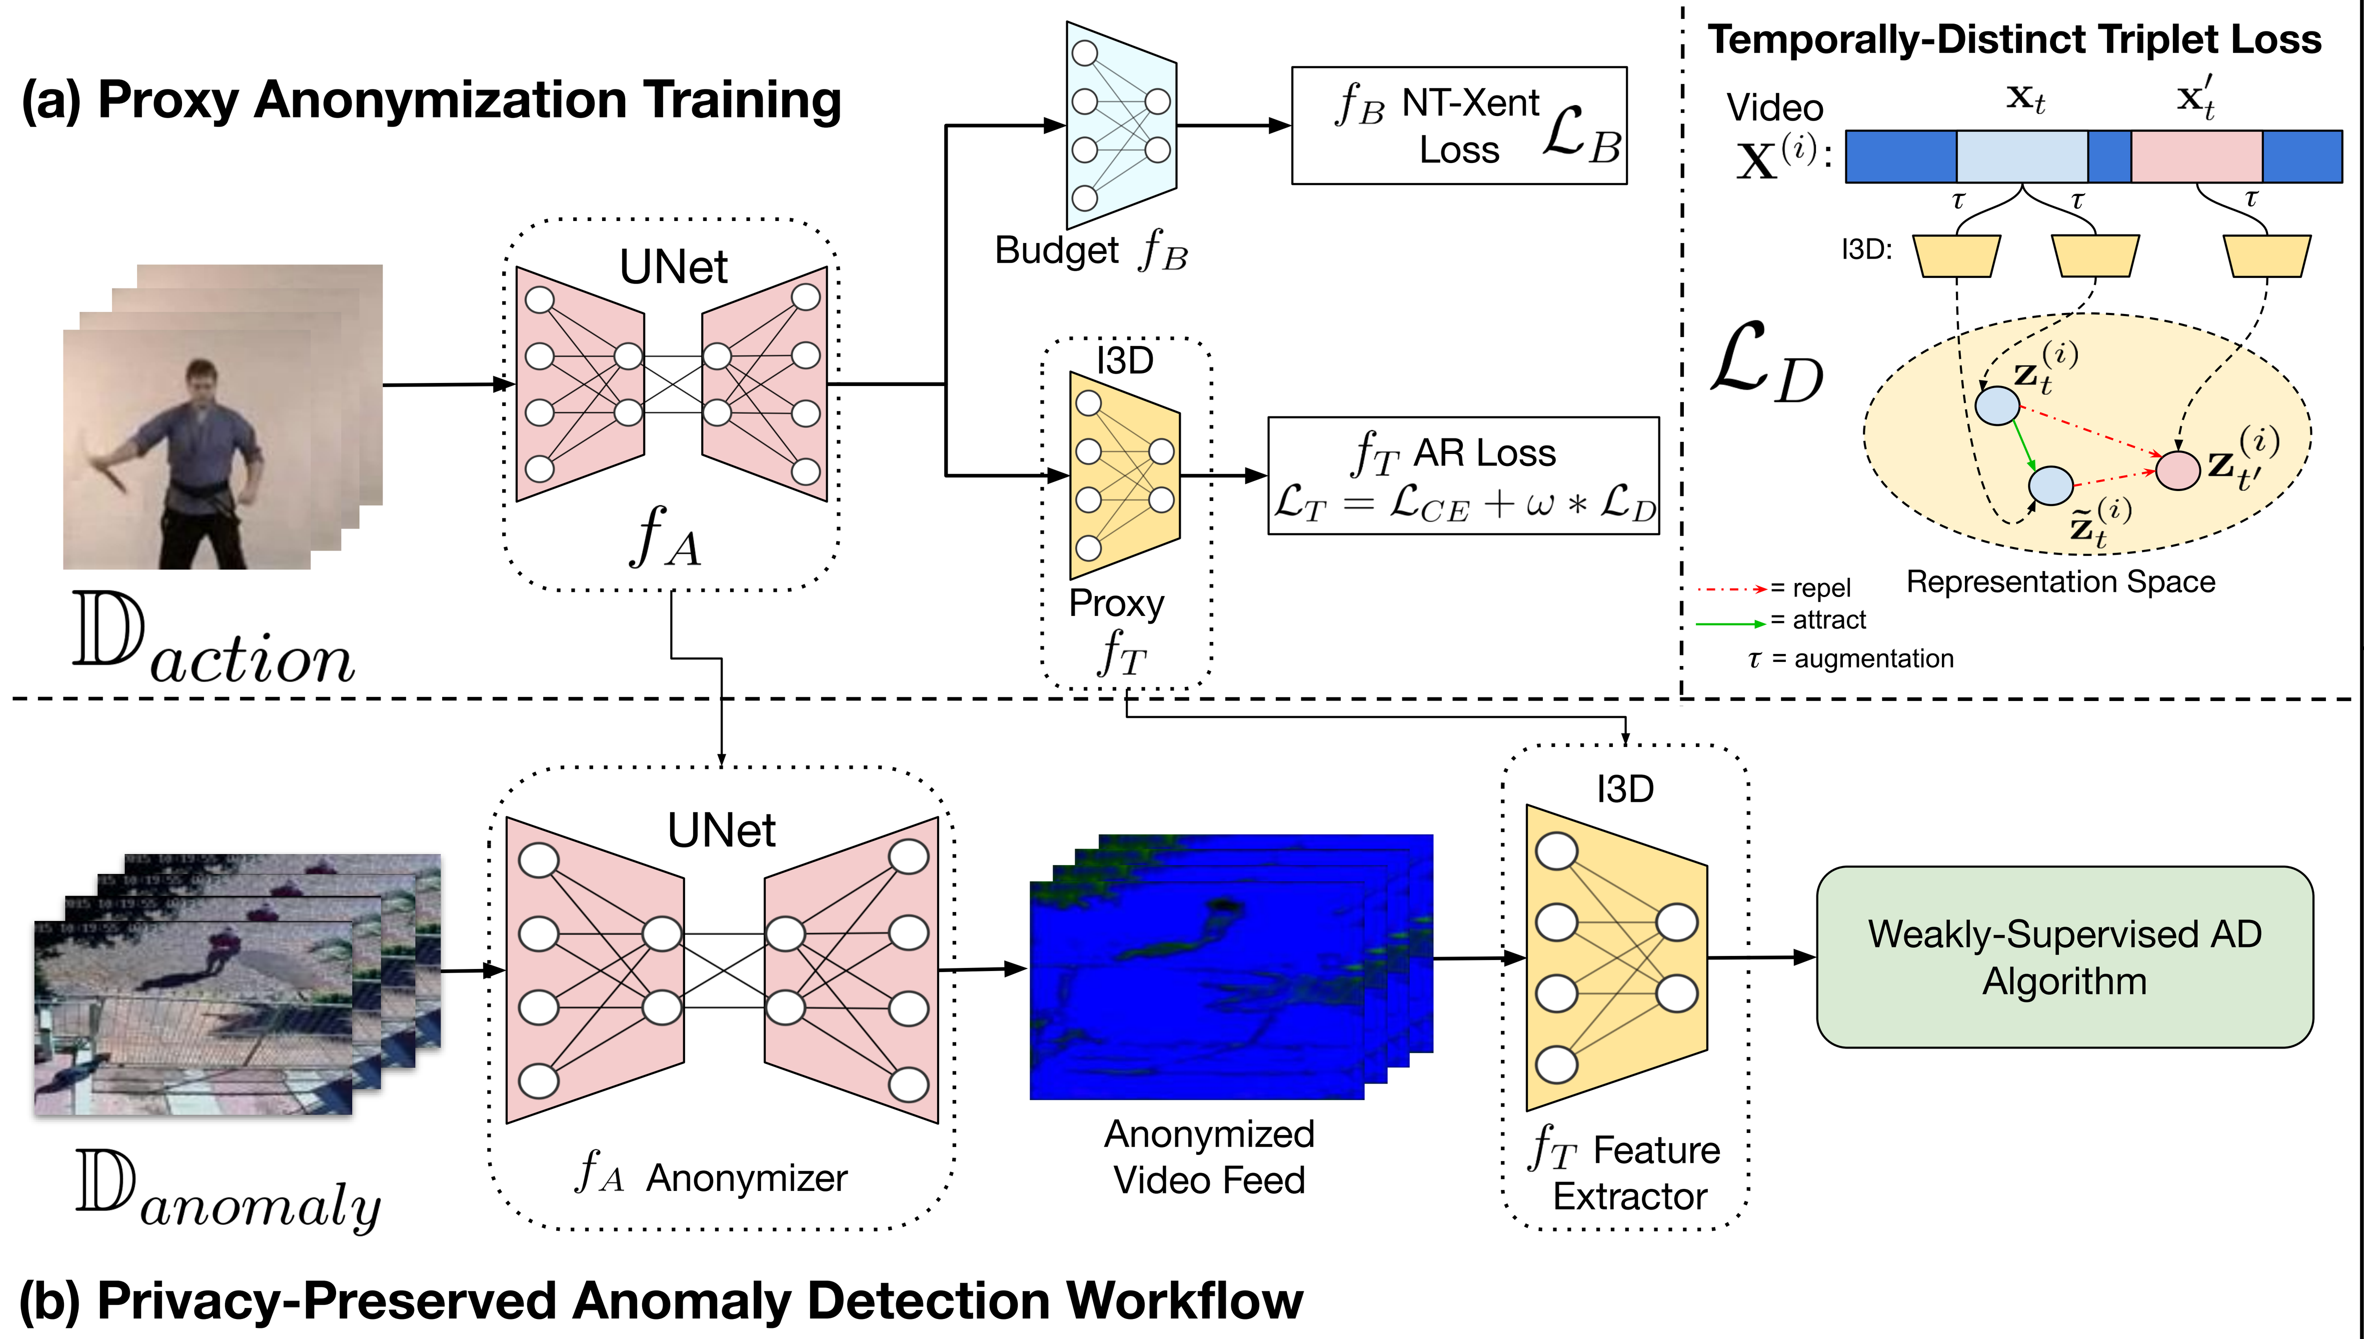

Comparison with existing privacy-preservation techniques

Method	VISPR Privacy cMAP(%)(↓)	UCF-Crime Anomaly AUC(%)(↑)	XD-Violence Anomaly AP(%)(↑)	ShanghaiTech Anomaly AUC(%)(↑)
Raw data	62.30	77.68	73.72	90.63
Downsample-2x	55.64 ↓10.69%	76.09 ↓2.05%	62.11 ↓15.75%	84.65 ↓6.60%
Downsample-4x	52.84 ↓15.18%	68.12 ↓12.31%	59.36 ↓19.48%	82.96 ↓8.46%
Obf-Blurring	58.68 ↓5.81%	75.69 ↓12.56%	56.17 ↓23.81%	89.63 ↓1.10%
Obf-Blackening	56.36 ↓9.53%	73.91 ↓4.85%	54.01 ↓26.74%	88.72 ↓2.11%
SPAct [1]	52.71 ↓15.39%	73.93 ↓4.83%	53.36 ↓27.62%	87.72 ↓3.21%
Ours	**42.21 ↓32.25%**	**74.81 ↓3.69%**	**60.32 ↓18.18%**	**90.59 ↓0.04%**

Results

Robbery:

Skateboarding:

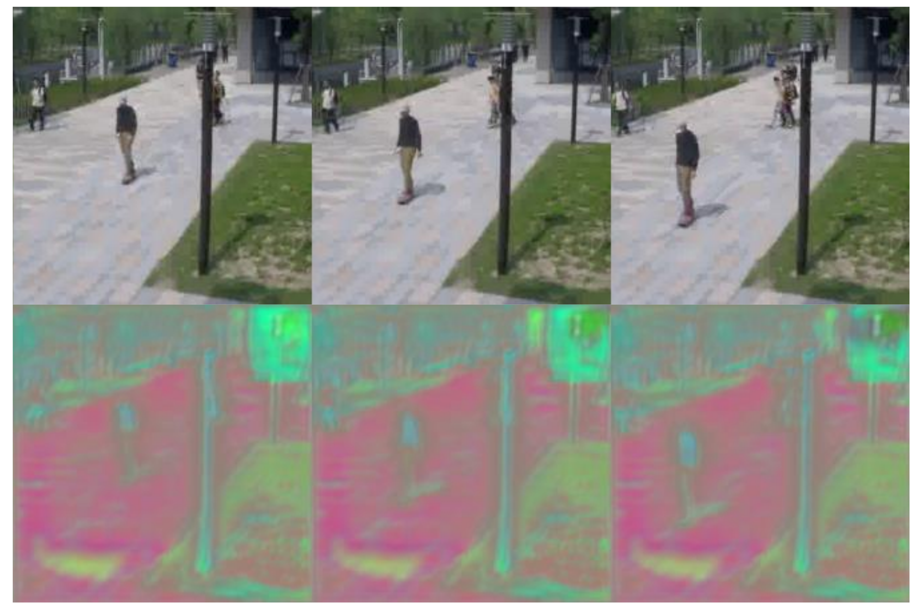

TeD-SPAD Training Algorithm

Algorithm 1: TeD-SPAD Framework

Inputs:
Datasets: $\mathbb{D}_{action}$, $\mathbb{D}_{anomaly}$
#Epochs: max_anon_epoch, max_ad_epoch
Learning Rates: α_{AD}, α_B, α_T
Hyperparameters: μ, ω
Output: θ_{AD}, θ_A
Model Initialization:
Initialize θ_T with Kinetics400 weights [4];
Initialize θ_B with SimCLR ImageNet weights [8].
Initialize $\theta_A \leftarrow \theta_A - \alpha_A \nabla_{\theta_A}(\mathcal{L}_{L1}(\theta_A))$(Ref. Eq. 3)
Anonymization Training:
for $e_0 \leftarrow 1$ **to** max_anon_epoch **do**
  Step-1
  $\theta_A \leftarrow \theta_A - \alpha_A \nabla_{\theta_A}(\mathcal{L}_T(\theta_A, \theta_T) - \omega L_B(\theta_A, \theta_B))$
  Step-2
  $\theta_T \leftarrow \theta_T - \alpha_T \nabla_{\theta_T}(\mathcal{L}_T(\theta_T, \theta_A))$, (Ref. Eq. 6)
  $\theta_B \leftarrow \theta_B - \alpha_B \nabla_{\theta_B}(\mathcal{L}_B(\theta_B, \theta_A))$.
end
Feature Extraction on $\mathbb{D}_{anomaly}$:
$\mathbb{F}_{anomaly} = \{ f_T(f_A(X^i)) \mid \forall X^i \in \mathbb{D}_{anomaly} \}$
Privacy-Preserved Anomaly Detection Training:
for $e_0 \leftarrow 1$ **to** max_ad_epoch **do**
  $\theta_{AD} \leftarrow \theta_{AD} - \alpha_{AD} \nabla_{\theta_{AD}}(L_{AD}(\theta_{AD}, \mathbb{F}_{anomaly}))$
end

Trade-off Plots