

Timestep-Compressed Attack on Spiking Neural Networks Through Timestep-Level Backpropagation

Donghwa Kang¹, Doohyun Kim¹, Sang-Ki Ko², Jinkyu Lee^{3,*}, Hyeongboo Baek^{2,†}, Brent ByungHoon Kang^{1,†}

¹Korea Advanced Institute of Science and Technology (KAIST)

²University of Seoul

³Yonsei University

hbbaek@uos.ac.kr, brentkang@kaist.ac.kr

Abstract

State-of-the-art (SOTA) gradient-based adversarial attacks on spiking neural networks (SNNs), which largely rely on extending FGSM and PGD frameworks, face a critical limitation: substantial attack latency from multi-timestep processing, rendering them infeasible for practical real-time applications. This inefficiency stems from their design as direct extensions of ANN paradigms, which fail to exploit key SNN properties. In this paper, we propose the timestep-compressed attack (TCA), a novel framework that significantly reduces attack latency. TCA introduces two components founded on key insights into SNN behavior. First, timestep-level backpropagation (TLBP) is based on our finding that global temporal information in backpropagation to generate perturbations is not critical for an attack’s success, enabling per-timestep evaluation for early stopping. Second, adversarial membrane potential reuse (A-MPR) is motivated by the observation that initial timesteps are inefficiently spent accumulating membrane potential, a warm-up phase that can be pre-calculated and reused. Our experiments on VGG-11 and ResNet-17 with the CIFAR-10/100 and CIFAR10-DVS datasets show that TCA significantly reduces the required attack latency by up to 56.6% and 57.1% compared to SOTA methods in white-box and black-box settings, respectively, while maintaining a comparable attack success rate.

1 Introduction

Spiking neural networks (SNNs) are often regarded as the third generation of artificial neural networks, known for their event-driven design and strong similarity to biological processes. Unlike conventional artificial neural networks (ANNs), SNNs use binary sequences to communicate between spiking neurons, producing spikes only when the membrane potential exceeds a certain threshold during the inference over multiple timesteps. This unique feature makes SNNs well-suited for handling spatio-temporal

*Jinkyu Lee initiated this work while affiliated with Sungkyunkwan University.

†Hyeongboo Baek and Brent ByungHoon Kang are the corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

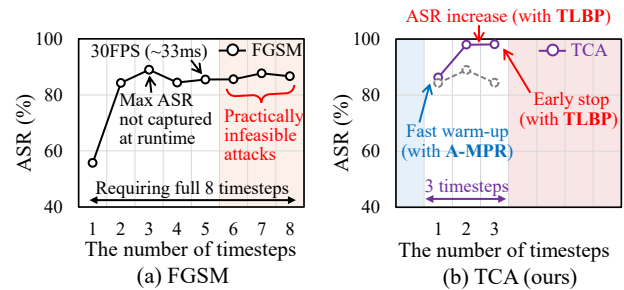


Figure 1: ASR over the number of timesteps for (a) conventional FGSM and (b) TCA, both utilizing STBP, on CIFAR-10 for VGG-11. TCA uses A-MPR for a fast warm-up and TLBP for early stopping to achieve high ASR with fewer timesteps compared to FGSM.

data and helps reduce power consumption (Tavanaei et al. 2019). Beyond image classification, SNNs have shown effectiveness in tasks such as text recognition and object detection. (Su et al. 2023; Zhu et al. 2023)

Despite their general resilience (Sharmin et al. 2020), the vulnerability of SNNs to adversarial attacks poses a significant security risk. This has led to the development of gradient-based attacks, which have evolved along two main paths. The first involves extending de facto standard attack methods (i.e., FGSM (Goodfellow, Shlens, and Szegedy 2014), PGD (Kurakin, Goodfellow, and Bengio 2018)) from ANNs to the SNN domain. The second focuses on developing SNN-specific gradient generation methods to accommodate their unique temporal dynamics. Key techniques include STBP (Wu et al. 2018), which adapts backpropagation through time (BPTT) for SNNs; RGA (Bu et al. 2023), which approximates gradients based on firing rates; and HART (Hao et al. 2023), a hybrid approach integrating both temporal and rate information. Crucially, these SNN-specific methods primarily serve to compute the gradients that are then fed into the established FGSM and PGD frameworks to craft the final attack.

A critical limitation of SNN attack methods exploiting

FGSM or PGD is the substantial latency from multi-timestep processing; for instance, the STBP-based attack in Fig. 1(a) requires eight full forward-backward passes to generate a single adversarial perturbation. While the specific “real-time” deadline depends on the application, a processing rate of 30 FPS is often considered the threshold for real-time performance in many fields, including object recognition in autonomous vehicles or CCTV systems (Zheng, Sun, and Dong 2025; Shi, Wang, and Guo 2024; Charles, Kazemipour, and Paninski 2024). Under such a widely used criterion, the latency of an eight-timestep attack becomes practically infeasible as it exceeds the operational deadline; PGD-based attacks are even slower. This suggests that SOTA attacks on SNNs remain largely theoretical, stemming from their design as direct extensions of ANN paradigms that fail to leverage two SNN-specific properties. First, there is a distinct latency-ASR (attack success ratio) trade-off (as shown in Fig. 1) which is also highly input-dependent; the number of timesteps needed to succeed varies significantly across inputs. This presents the challenge of developing a runtime mechanism to capture this input-dependent vulnerability for early stopping. Second, this trade-off is non-linear, as ASR often increases sharply in the initial timesteps. This poses a further challenge of how to selectively allocate computational resources across different attack stages to minimize overall latency.

In this paper, we propose a novel timestep-compressed attack (TCA) that establishes such a new paradigm, circumventing the reliance on full-timestep information to significantly reduce the required number of timesteps, illustrated in Fig. 1(b). TCA incorporates two primary components targeting a classification problem. First, timestep-level backpropagation (TLBP) employs a per-timestep strategy for generating perturbations (while eliminating BPTT) to the input image. This approach enables cross-entropy monitoring at each timestep for early stopping and, counter-intuitively, significantly improves the final ASR compared to the existing FGSM and PGD using full-timestep temporal information. Second, adversarial membrane potential reuse (A-MPR) facilitates a rapid warm-up by initializing the attack with pre-calculated membrane potentials. This “membrane image” is generated offline via a purpose-built loss function that balances gradient accuracy with adversarial efficacy to ensure the extraction of effective gradients from the onset of the attack.

We implemented TCA to all STBP, RGA, and HART, and it supports both white-box and black-box attacks. In line with the SOTA studies, we conducted experiments on CIFAR-10/100 and CIFAR10-DVS for VGG-11 and ResNet-17. Our experiments validate TCA’s effectiveness on VGG-11 and ResNet-17 across the CIFAR-10 and CIFAR-100 datasets. While preserving a comparable ASR, TCA lowers the required attack latency by up to 56.6% in white-box and 57.1% in black-box scenarios, respectively, against SOTA methods.

The main contributions of this paper are as follows:

- We analyze the impact of global temporal information within the backpropagation process of gradient-based attacks, revealing its minimal effect on the final ASR.

- Based on this analysis, we propose TLBP, a novel per-window backpropagation method that enables input-dependent early stopping for low-latency attacks.
- We propose A-MPR, an offline, rapid warm-up technique that pre-calculates and reuses adversarial membrane potentials to overcome the inefficiency of early attack timesteps.
- We demonstrate the effectiveness of our framework through extensive experiments, showing a superior latency-ASR trade-off compared to SOTA methods in both white-box and black-box scenarios.

2 Related Work

Learning method of SNNs. The two primary methods for training SNNs are ANN-SNN conversion (Cao, Chen, and Khosla 2015) and direct training (Wu et al. 2018). ANN-SNN conversion involves training a conventional ANN and then mapping its weights to an SNN of the same architecture. While converted SNNs can achieve comparable accuracy to ANNs without retraining, they typically require a large number of timesteps for inference, though recent works have sought to mitigate this (Deng and Gu 2021; Diehl et al. 2015). Alternatively, direct training uses methods like STBP to train the SNN from scratch. STBP was proposed by Wu et al. (Wu et al. 2018) who recognized the similarity between the temporal dynamics of SNNs and RNNs, adapting the BPTT (Werbos 2002) algorithm. Direct training with STBP enables high accuracy with fewer timesteps than conversion methods, but traditionally struggled to match the peak accuracy of ANNs. However, recent studies have made progress in closing this accuracy gap (Fang et al. 2021; Lee, Delbruck, and Pfeiffer 2016; Neftci, Mostafa, and Zenke 2019).

Adversarial attacks for SNNs. The vulnerability of ANNs to adversarial examples—subtle, imperceptible input perturbations that cause misclassification—was first demonstrated by Szegedy et al. (Szegedy et al. 2013). This led to the development of gradient-based attack methods such as FGSM (Goodfellow, Shlens, and Szegedy 2014) and PGD (Kurakin, Goodfellow, and Bengio 2018), which generate perturbations using the model’s gradients. These methods have since been extended to SNNs, typically using the STBP gradient to guide the perturbation (Sharmin et al. 2019). However, the non-differentiable nature of spiking neurons makes direct gradient calculation inefficient. To overcome this, alternative backpropagation methods have been developed that approximate gradients using firing rates or other temporal information. For example, RGA (Bu et al. 2023) approximates gradients based on firing ratios, while HART (Hao et al. 2023) uses a hybrid of firing ratio and temporal information. A common drawback of these approaches, however, is that their reliance on the FGSM and PGD frameworks necessitates performing a forward and/or backward pass for each timestep in the sequence.

3 Preliminaries

In this section, we present the foundational mechanisms of SNNs and outline the principles of adversarial attacks. As

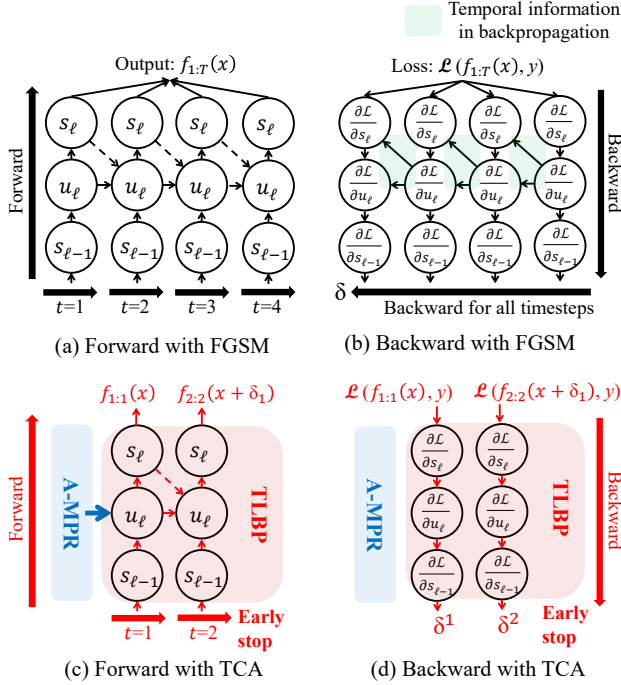


Figure 2: Forward pass and backward pass for FGSM (a–b) and TCA (c–d) for window size $w = 1$. Since TCA does not rely on temporal information between consecutive windows, it can decide whether to stop the attack at each window using TLBP and measured cross-entropy. The blue line indicates the reuse of pre-extracted membrane potentials for each class through A-MPR for rapid warm-up.

illustrated in Fig. 2(a), an SNN propagates an input image (denoted by x) across multiple timesteps and layers. We formulate this forward pass, denoted by $f_{1:T}$, as follows:

$$f_{1:T}(x) = \frac{1}{T} \sum_{t=1}^T h \circ l^L \circ l^{L-1} \circ \dots \circ l^1(x), \quad (1)$$

where T and L are the number of timesteps and blocks, respectively. $l^{\ell}(\cdot)$ represents the ℓ -th block, and each block consists of a convolution layer, a batch normalization layer, and a leaky integrated-and-fire (LIF) layer (Gerstner et al. 2014). $h(\cdot)$ denotes the classifier in the last layer of the network.

At each timestep, LIF neurons accumulate membrane potential across T . The membrane potential $u_{\ell,t}$ at the t -th timestep in ℓ -th block is given by:

$$u_{\ell,t} = \tau u_{\ell,t-1} + W_{\ell} s_{\ell-1,t} + \eta_{\ell,t}, \quad (2)$$

where τ and W_{ℓ} denote the leaky factor and the convolutional weight in the ℓ -th block, respectively (Brette et al. 2007). $\eta_{\ell,t}$ denotes the reset factor that is applied to the membrane potential $u_{\ell,t}$ when it exceeds a certain threshold. $s_{\ell-1,t}$ represents the spike fired at the t -th timestep in $(\ell - 1)$ -th block. We formulate $s_{\ell-1,t}$ as follows:

$$s_{\ell-1,t} = \begin{cases} 1 & \text{if } u_{\ell-1,t} > V_{th}, \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

where V_{th} is predefined threshold. If $u_{\ell-1,t}$ exceeds the threshold V_{th} , the LIF neuron fires a spike to the next layer and resets its membrane potential.

Adversarial attacks are approaches used to deceive SNNs into making incorrect predictions by generating imperceptible perturbations. The formula to generate perturbations (illustrated in Fig. 2(b)) is defined as the following optimization problem:

$$\delta = \arg \max_{\delta, \|\delta\|_{\infty} \leq \epsilon} \mathcal{L}(f_{1:T}(x + \delta|\theta), y), \quad (4)$$

where δ is a generated perturbation and \mathcal{L} is a loss function. ϵ is a constant value that ensures the perturbation remains imperceptible. $f_{1:T}(\cdot|\theta)$ is an SNN model parameterized by θ , and y is the target label intended to mislead.

We consider two basic gradient-based attack methods applied in both scenarios: FGSM (Goodfellow, Shlens, and Szegedy 2014) and PGD (Kurakin, Goodfellow, and Bengio 2018). FGSM is an attack method that generates a perturbation by calculating the sign of the gradients against the loss function, that is

$$\delta = \epsilon \cdot \text{sgn}(\nabla_{\delta} \mathcal{L}(f_{1:T}(x|\theta), y)), \quad (5)$$

where $\text{sgn}(\cdot)$ represents the function to extract sign of the gradients. PGD, an advanced extension of FGSM, generates refined perturbations through iterative gradient extraction (Madry et al. 2017). PGD adjusts the size of perturbations γ for the given total iterations, which is expressed as follows:

$$\delta^r = \text{clip}_{[-\epsilon, \epsilon]}(\delta^{r-1} + \gamma \cdot \text{sgn}(\nabla_{\delta^{r-1}} \mathcal{L}(f_{1:T}(x + \delta^{r-1}|\theta), y))), \quad (6)$$

where r is the index of iteration (for $1 \leq r$, $\delta^0 = 0$), and $\text{clip}(\cdot)$ ensures that the perturbation satisfies $\|\delta\|_{\infty} \leq \epsilon$. However, applying these attack techniques to SNNs requires at least T forward and backward propagations.

4 TCA: Methodology

This section details TCA, a method targeting classification problems in both white-box and black-box settings, motivated by real-time applications (e.g., traffic sign recognition in autonomous vehicles). Our threat model is grounded in practical feasibility. In a scenario without any system access, the attack can be physically realized by projecting optical perturbations onto objects (e.g., traffic signs), a technique demonstrated by AttackZone (Muller et al. 2022). Alternatively, with limited physical access such as USB attachment, a man-in-the-middle (MITM) attack—where the image stream is intercepted and reinjected after perturbation—becomes feasible, as shown in recent studies (Wu, Rowlands, and Wahlström 2024; Liu et al. 2022).

4.1 Observations

We present two measurement-based observations that reveal the inefficiencies of conventional SNN attacks and form the foundation for our proposed method, TCA.

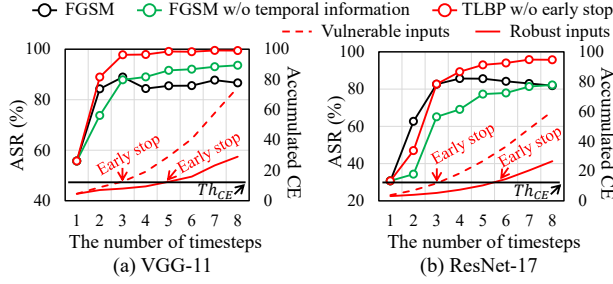


Figure 3: ASR and accumulated CE (calculated by $\mathcal{L}(f_{t_s:t_e}(\cdot), y)$) over timesteps on CIFAR-10 for three STBP variants (FGSM, FGSM w/o temporal information, and TLBP (w/o early stop, $w = 1$)), evaluated on (a) VGG-11 and (b) ResNet-17. Vulnerable and robust inputs refer to samples successfully attacked by TLBP at the first and eighth timesteps, respectively.

Necessity of temporal information. Our first analysis questions the reliance on full-timestep processing in conventional SNN attacks. As shown by the standard STBP-based FGSM attack in Fig. 3 (black line), the ASR quickly saturates, suggesting that processing all timesteps is often redundant. This standard approach requires a full forward pass followed by an extensive backward pass because the membrane potential at each timestep t is dependent on the state at $t - 1$ (Eq. (2)), analogous to BPTT in RNNs¹.

To investigate the necessity of this temporal dependency, we conducted a motivational experiment with an attack variant that disregards it (i.e., eliminating BPTT). This approach draws inspiration from an emerging theme in SOTA research, where shortening the temporal backpropagation path is used to mitigate challenges in direct SNN training, such as vanishing gradients. Our variant calculates gradients on a per-timestep basis, effectively treating each timestep as independent. Despite sacrificing this global temporal information, the resulting attack (green line in Fig. 3) achieves a comparable final ASR to the standard method. This finding leads to our first key observation:

- O1. The global temporal information aggregated across all timesteps has a minimal impact on the final ASR of gradient-based attacks.

Inefficacy of early timesteps. The second challenge arises from the initial timesteps of an attack. Fig. 4 shows that both the neuron firing rate and the corresponding ASR of FGSM (black line) are negligible at the beginning of the attack sequence. This indicates that insufficient information propagates through the network in the early stages, resulting in ineffective gradients. This forms our second observation:

- O2. Gradients extracted from the initial timesteps are ineffective due to low neuron firing activity, hindering the attack process from the start.

¹Note that we consider the direct training method for SNNs, which adapts the BPTT algorithm by using a surrogate gradient to handle the non-differentiable LIF activation function.

4.2 TLBP: Timestep-Level Backpropagation

Building upon O1, TLBP enables input-dependent early termination by monitoring an attack’s success at runtime, a capability absent in conventional FGSM and PGD. Instead of processing all timesteps at once, TLBP operates on a “window,” a unit of w timesteps ($1 \leq w \leq T$) over which a single forward-backward pass is performed to generate a perturbation. This windowed approach generalizes the conventional attacks; if the window size w is set to the total number of timesteps T , a single iteration of TLBP is equivalent to FGSM, while multiple iterations are equivalent to PGD.

The operational flow of TLBP is illustrated in Fig. 2(c) and (d) for a window size of $w = 1$. In n -th window, the forward pass (Fig. 2(c)) computes the SNN output over the current window, $f_{t_s:t_e}(x + \delta^{n-1})$. This output is used to calculate the CE loss, $\mathcal{L}(f_{t_s:t_e}(\cdot), y)$, which serves as the input to the backward pass (Fig. 2(d)). The backward pass then computes the gradient to update the perturbation from δ^{n-1} to δ^n as follows:

$$\delta^n = \text{clip}_{[-\epsilon, \epsilon]}(\delta^{n-1} + \epsilon \cdot \text{sgn}(\nabla_{\delta^{n-1}} \mathcal{L}(f_{t_s:t_e}(x + \delta^{n-1}|\theta), y))), \quad (7)$$

where $t_s = (n - 1) \cdot w + 1$ and $t_e = \min(n \cdot w, T)$ define the window boundaries for the n -th window ($1 \leq n, \delta^0 = 0$). This per-window strategy not only provides a robust criterion for early stopping but also counter-intuitively improves the final ASR compared to full-timestep methods, as evidenced by the significant performance gap between TLBP (red line with circle markers) and standard FGSM (black line) in Fig. 3.

To determine when to stop, TLBP uses cross-entropy (CE) as a proxy for attack success. A high CE loss signifies that the model’s confidence in the true class is low, indicating that the perturbation is successfully steering the prediction toward the adversarial target. The attack terminates once the accumulated CE, defined as $\sum_{n=1} \mathcal{L}(f_{t_s:t_e}(\cdot), y)$, reaches a predefined threshold Th_{CE} (or when $t_s \geq T$ holds). The validity of this metric is empirically demonstrated in Fig. 3, where inputs vulnerable to attacks exhibit a correspondingly rapid accumulation of CE.

4.3 A-MPR: Adversarial Membrane Potential Reuse

A-MPR is a rapid warm-up technique that addresses the inefficiency of early attack timesteps based on O2. It overcomes the “cold start” problem of SNNs by initializing the network with pre-calculated, non-zero membrane potentials, ensuring effective gradient extraction from the onset of the attack. The insight for A-MPR comes from observing the negligible initial spike ratios in conventional attacks (Fig. 4(a, c)). This indicates a wasteful warm-up phase where neurons only accumulate potential without firing, yielding less useful gradients. A-MPR bypasses this inefficient runtime phase by pre-calculating and reusing these initial membrane potentials.

Since randomly generating these potentials can degrade ASR, we pre-generate “membrane image” y^* for each attack target y using a purpose-built loss function that balances

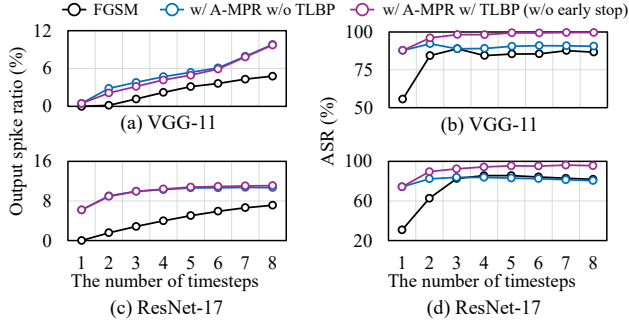


Figure 4: Output spike ratio (a, c) and ASR (b, d) for three STBP-based methods—a standard FGSM, an attack with A-MPR, and one combining A-MPR with TLBP (without early stopping, $w = 1$)—evaluated over different timesteps on VGG-11 and ResNet-17 with CIFAR-10.

two objectives: (i) ensuring the generated potentials are similar to those from real inputs for gradient accuracy, and (ii) steering the initial inference away from the correct label for adversarial efficacy. We propose \mathcal{L}_{mem} for objective (i) and \mathcal{L}_{adv} for (ii). The similarity loss, \mathcal{L}_{mem} , is formulated using cosine similarity ($\text{sim}(\cdot, \cdot)$) as follows:

$$\mathcal{L}_{mem}(t_1) = \sum_{\ell=1}^L \text{sim}(u_{\ell, t_1}, u_{\ell, t_1}^*), \quad (8)$$

where u_{ℓ, t_1} and u_{ℓ, t_1}^* are the membrane potentials at timestep t_1 and block ℓ accumulated by the input image x and the trainable membrane image y^* , respectively. The adversarial loss, \mathcal{L}_{adv} , encourages the pre-calculated potentials to already be adversarial:

$$\mathcal{L}_{adv}(x, y, t_1) = \mathcal{L}(f_{t_1+1:T}(x), y). \quad (9)$$

The total loss for generating the membrane image is the sum of these two components:

$$\mathcal{L}_{mpr}(x, y, t_1) = \mathcal{L}_{mem}(t_1) + \mathcal{L}_{adv}(x, y, t_1). \quad (10)$$

The membrane image is then optimized via gradient ascent for each class in the training set:

$$y^* = y^* + \beta \cdot \text{sgn}(\nabla_{y^*} \mathcal{L}(x, y, t_1)), \quad (11)$$

where β is the update parameter. Before an attack, the network is warmed up by reusing the prepared y^* for the target label y . As a post-processing step, we clamp all membrane potentials to a range of $[V_{min}, V_{max}]$ to mitigate the effect of outliers. Since the entire A-MPR process is performed offline, it introduces no computational overhead to the runtime attack.

The results in Fig. 4(b, d) validate this approach. A-MPR alone (blue line) significantly boosts the initial ASR over the standard FGSM (black line). When combined with TLBP (purple line), the ASR saturates even more rapidly, creating an ideal condition for the early stopping criterion from Section 4.2 to effectively reduce latency. The comprehensive latency-ASR performance of the complete TCA framework is demonstrated in the evaluation section.

5 Evaluation

In this section, we evaluate the effectiveness of TCA against existing attack methods in both white-box and black-box scenarios. Our experiments are conducted on VGG-11 (Simonyan and Zisserman 2014) and ResNet-17 (Zheng et al. 2021) architectures with the CIFAR-10/100 (Krizhevsky, Hinton et al. 2009) and CIFAR10-DVS (Simonyan and Zisserman 2014) datasets². All experiments are conducted on NVIDIA RTX A6000 GPUs. All victim SNNs are trained using the STBP method, as ANN-SNN converted models have known security vulnerabilities. We compare the performance of our proposed TCA framework against the baseline FGSM and PGD attacks, evaluating each in combination with three SOTA gradient generation methods: STBP, RGA, and HART. Performance is measured using two metrics: ASR, the proportion of samples that cause a misclassification, and the attack latency, the total time required to generate a perturbation for a single input.

5.1 Comparison to Prior Art

We evaluate the latency-ASR effectiveness of TCA against FGSM and PGD baselines in Figs. 5 and 6. The experiments are conducted on VGG-11 and ResNet-17 with CIFAR-10/100 datasets in an untargeted scenario, using $\epsilon = 8/255$, $t_1 = 2$, $\beta = 8/255$, $V_{min} = 0.2$, and $V_{max} = 0.5$. To ensure a fair comparison at low latencies, PGD is restricted to two iterations with a step size of $8/255$. The points for the baselines represent attacks with a fixed number of timesteps from one to eight. In contrast, TCA operates flexibly; we plot two execution options. The hollow red points represent attacks using only the first window ($n = 1$) with varying window sizes (w). The solid red points represent attacks with a fixed w and varying Th_{CE} values (0.01, 0.1, 1, or 5), which determines the number of iterations (windows) n . Based on our motivation in Sec. 1, we consider attacks that complete within 33ms (equivalent to 30 FPS) to be practically feasible for real-time systems.

White-box scenario. As shown in Fig. 5, most PGD configurations are practically infeasible (in the red square area of each subfigure) due to their iterative nature, while FGSM, though feasible, achieves a relatively low ASR. For TCA, the results (solid red points) show that Th_{CE} acts as a control-knob for the latency-ASR trade-off; a lower threshold allows the attack to stop at an earlier window (e.g., $n = 1$), while a higher threshold requires more windows (e.g., $n = 4$). TCA demonstrates a superior latency-ASR trade-off compared to both FGSM and PGD across most settings (e.g., up to 56.6% latency reduction for STBP on VGG-11 and CIFAR-10), while its performance on the VGG-11 and CIFAR-100 combination remains comparable.

Black-box scenario. In the black-box scenario, we assume the attacker knows only the model architecture, not any specific parameters (following the same black-box assumptions of RGA and HART for a fair comparison). Using the same as the white-box setting (e.g., ϵ , t_1 , \dots , w , n),

²Results for CIFAR10-DVS are provided in the supplement due to page limits.

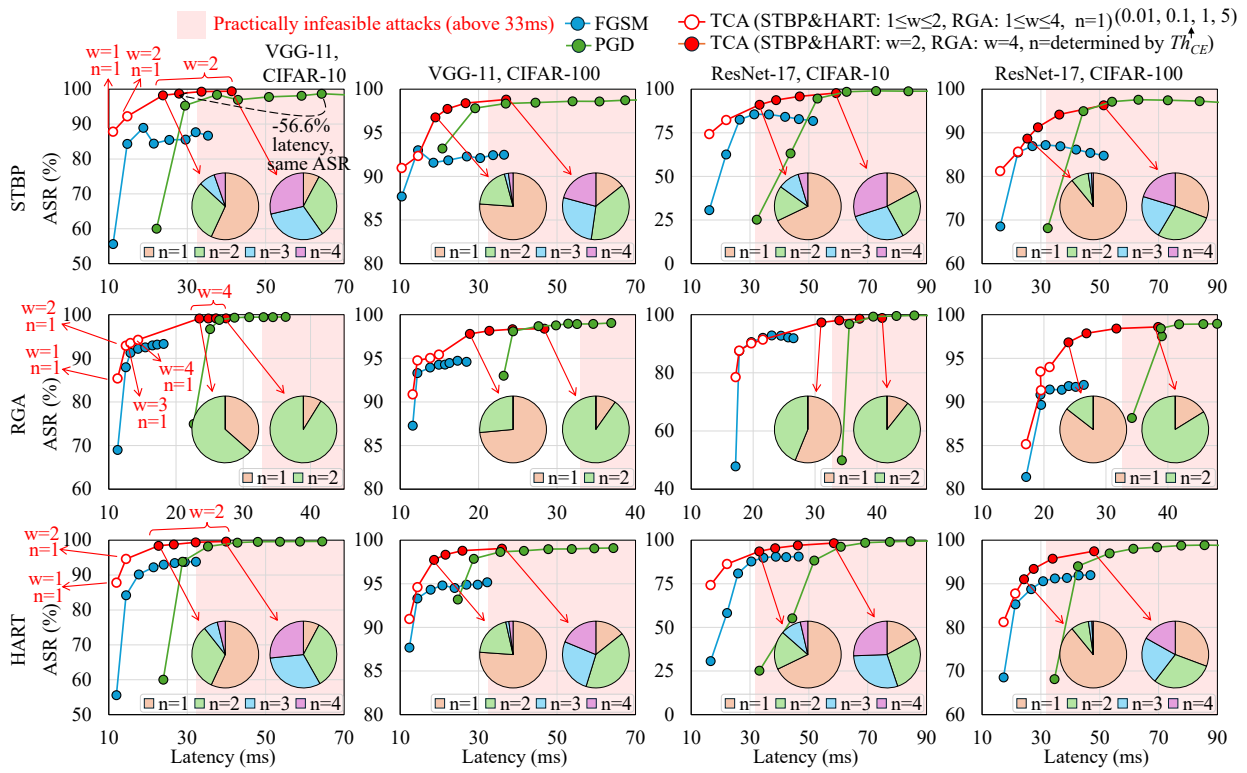


Figure 5: Experimental results on whitebox scenario

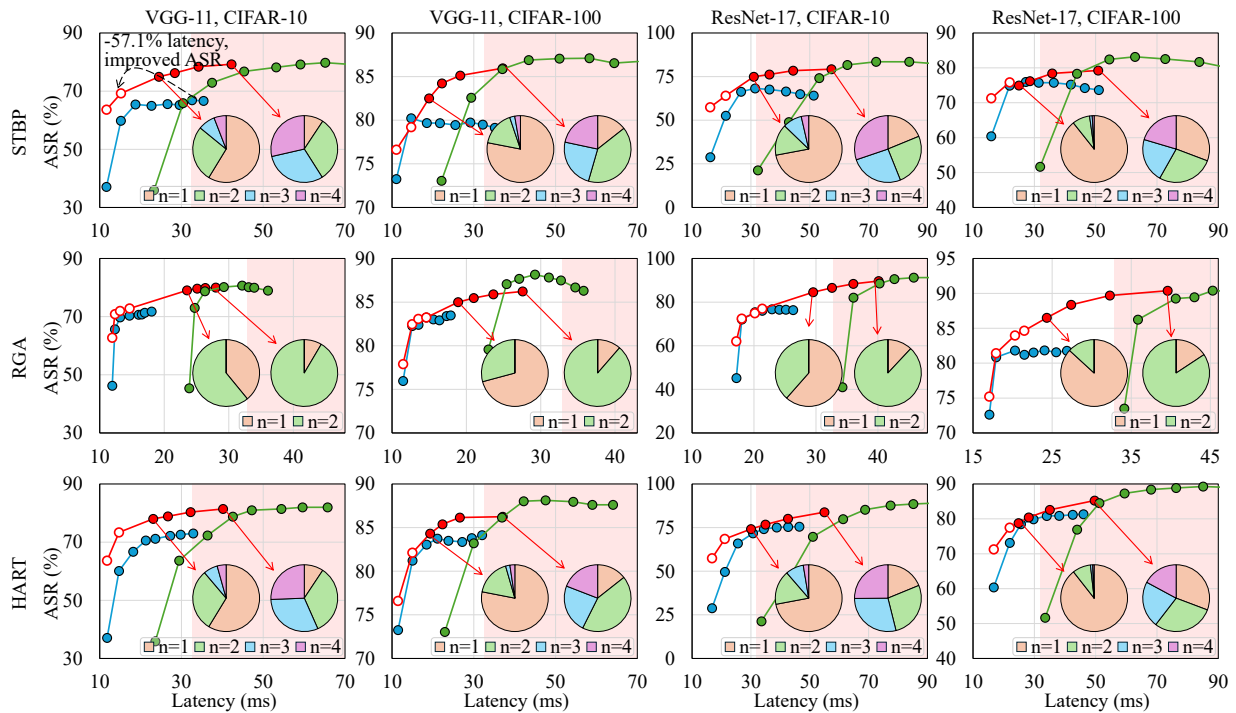


Figure 6: Experimental results on blackbox scenario

BP	\mathcal{L}_{adv}	\mathcal{L}_{mem}	y -specific	VGG-11		ResNet-17	
				$w=1$	$w=2$	$w=1$	$w=2$
	✓	✗	✗	84.7	89.1	74.2	80.4
STBP	✓	✓	✗	85.6	88.0	74.8	80.6
	✓	✓	✓	87.8	92.2	74.3	82.3
	✓	✗	✗	82.1	91.4	76.9	85.6
RGA	✓	✓	✗	84.4	92.2	78.2	85.1
	✓	✓	✓	85.4	92.9	78.5	87.6
	✓	✗	✗	84.0	92.5	73.0	84.0
HART	✓	✓	✗	85.2	92.4	74.5	85.0
	✓	✓	✓	87.8	94.6	74.3	86.3

Table 1: Ablation study of A-MPR on CIFAR-10

BP	(V_{min}, V_{max})	VGG-11		ResNet-17	
		$w=1$	$w=2$	$w=1$	$w=2$
	(none, none)	14.8	17.0	25.4	29.8
	(0.0, 1.0)	60.8	82.7	36.8	55.8
STBP	(0.1, 0.8)	78.8	91.8	53.6	72.0
	(0.2, 0.6)	89.6	93.5	71.7	80.6
	(0.2, 0.5)	87.8	92.2	74.3	82.3
	(none, none)	29.1	37.9	36.8	47.5
	(0.0, 1.0)	71.3	91.2	51.4	77.4
RGA	(0.1, 0.8)	79.0	93.7	65.1	85.2
	(0.2, 0.6)	87.4	93.9	76.4	88.1
	(0.2, 0.5)	85.4	92.9	78.5	87.6
	(none, none)	14.8	19.0	25.4	29.6
	(0.0, 1.0)	60.8	84.6	36.6	56.5
HART	(0.1, 0.8)	78.8	92.8	53.2	75.8
	(0.2, 0.6)	89.6	95.7	71.6	85.1
	(0.2, 0.5)	87.8	94.6	74.3	86.3


Table 2: Various minimum values of membrane potentials on CIFAR-10

TCA again demonstrates a superior latency-ASR trade-off compared to the baselines across most configurations (e.g., up to 57.1% latency reduction for STBP on VGG-11 and CIFAR-10), as shown in Fig. 6, despite a marginally lower ASR on the VGG-11 and CIFAR-100 combination.

5.2 Ablation Study

Impact of loss function in A-MPR. We ablate the components of our proposed A-MPR—the adversarial loss (\mathcal{L}_{adv}), the similarity loss (\mathcal{L}_{mem}), and the class (y)-specific strategy—to analyze their individual impact on TCA’s performance. The results are shown in Table 1, which presents the ASR for VGG-11 and ResNet-17 on CIFAR-10 under STBP, RGA, and HART attacks. To isolate the effect of A-MPR, we fix the number of iterations to one ($n = 1$) and test for window sizes $w = 1$ and $w = 2$. The results indicate that incrementally adding each component generally increases ASR, regardless of the model or backpropagation method. For instance, for VGG-11 with STBP at $w = 1$, the ASR increases from 84.7% to 87.8% with all components enabled.

Boundary of membrane potentials. We investigate the impact of clamping the membrane potentials to a range of $[V_{min}, V_{max}]$ during the A-MPR process. As shown in Table 2, not bounding the potentials (none, none) causes a sig-



(a) Vulnerable samples : hard to classify but easy to attack
(b) Robust samples : easy to classify but hard to attack

Figure 7: Visualization of vulnerable (attacked at $n = 1$) and robust (attacked at $n = 8$) samples on VGG-11 with CIFAR-100

nificant drop in ASR across all configurations, which we attribute to incorrectly accumulated potentials in outlier neurons. In contrast, setting the boundaries (e.g., $V_{min} = 0.2$, $V_{max} = 0.5$) substantially improves performance. For example, for ResNet-17 with STBP ($w = 1$, $n = 1$), the ASR increases from a very low 25.4% to as high as 74.3% when the potentials are appropriately bounded.

Attack vulnerability. Fig. 7 visualizes the characteristics of inputs from CIFAR-100 that are either vulnerable or robust to our attack on the VGG-11 model. We define ‘vulnerable’ samples as those successfully attacked in the first window ($w = 1$, $n = 1$), and ‘robust’ samples as those that resist the initial attack but are compromised in the 8-th window ($w = 1$, $n = 8$). The visualization indicates that vulnerable samples are often visually ambiguous (e.g., small objects or poor contrast), corresponding to low initial confidence scores for the ground truth. Consequently, these samples require only a small perturbation over a few timesteps to be misclassified, whereas robust samples, with their high-confidence visual clarity, demand a larger, cumulative perturbation.

6 Conclusion

In this paper, we addressed the critical limitation of high latency in SOTA adversarial attacks on SNNs, which renders them practically infeasible for real-time systems. We introduced TCA, a framework built on two key insights into SNN attack dynamics. We first analyzed the global temporal information in the backpropagation of the SNN attack and revealed its minimal impact on an attack’s final success. Based on this finding, we developed TLBP to perform per-window evaluation, enabling effective early stopping without sacrificing performance. To accelerate the initial attack phase, A-MPR bypasses the inefficient warm-up period by reusing pre-calculated potentials. Our experiments on VGG-11 and ResNet-17 with CIFAR-10/100 and CIFAR10-DVS validate this approach, demonstrating a reduction in required attack latency of up to 56.6% and 57.1% for white-box and black-box attacks, respectively, all while maintaining a comparable ASR to SOTA methods.

Acknowledgments

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT)(RS-2025-00520023, RS-2024-00438248, RS-2025-25424702). This research was also supported by the MSIT, Korea under the ITRC(Information Technology Research Center) support program(IITP-2025-RS-2023-00259061) supervised by the IITP(Institute for Information & Communications Technology Planning & Evaluation). This work was supported by a research project with NYU (G01240629).

References

- Brette, R.; Rudolph, M.; Carnevale, T.; Hines, M.; Beeman, D.; Bower, J. M.; Diesmann, M.; Morrison, A.; Goodman, P. H.; Harris Jr, F. C.; et al. 2007. Simulation of networks of spiking neurons: a review of tools and strategies. *Journal of computational neuroscience*, 23(3): 349–398.
- Bu, T.; Ding, J.; Hao, Z.; and Yu, Z. 2023. Rate gradient approximation attack threatens deep spiking neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7896–7906.
- Cao, Y.; Chen, Y.; and Khosla, D. 2015. Spiking deep convolutional neural networks for energy-efficient object recognition. *International Journal of Computer Vision*, 113(1): 54–66.
- Charles, A. S.; Kazemipour, A.; and Paninski, L. 2024. realSEUDO for real-time calcium imaging analysis. In *Advances in Neural Information Processing Systems* 37.
- Deng, S.; and Gu, S. 2021. Optimal conversion of conventional artificial neural networks to spiking neural networks. *arXiv preprint arXiv:2103.00476*.
- Diehl, P. U.; Neil, D.; Binas, J.; Cook, M.; Liu, S.-C.; and Pfeiffer, M. 2015. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *2015 International joint conference on neural networks (IJCNN)*, 1–8. IEEE.
- Fang, W.; Yu, Z.; Chen, Y.; Huang, T.; Masquelier, T.; and Tian, Y. 2021. Deep residual learning in spiking neural networks. *Advances in Neural Information Processing Systems*, 34: 21056–21069.
- Gerstner, W.; Kistler, W. M.; Naud, R.; and Paninski, L. 2014. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Hao, Z.; Bu, T.; Shi, X.; Huang, Z.; Yu, Z.; and Huang, T. 2023. Threaten spiking neural networks through combining rate and temporal information. In *The Twelfth International Conference on Learning Representations*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2018. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, 99–112. Chapman and Hall/CRC.
- Lee, J. H.; Delbruck, T.; and Pfeiffer, M. 2016. Training deep spiking neural networks using backpropagation. *Frontiers in neuroscience*, 10: 508.
- Liu, W.; He, W.; Hu, B.; and Chang, C.-H. 2022. A practical man-in-the-middle attack on deep learning edge device by sparse light strip injection into camera data lane. In *2022 IEEE 35th International System-on-Chip Conference (SOCC)*, 1–6. IEEE.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Muller, R.; Man, Y.; Celik, Z. B.; Li, M.; and Gerdes, R. 2022. Physical hijacking attacks against object trackers. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2309–2322.
- Neftci, E. O.; Mostafa, H.; and Zenke, F. 2019. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6): 51–63.
- Sharmin, S.; Panda, P.; Sarwar, S. S.; Lee, C.; Ponghiran, W.; and Roy, K. 2019. A comprehensive analysis on adversarial robustness of spiking neural networks. In *2019 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Sharmin, S.; Rathi, N.; Panda, P.; and Roy, K. 2020. Inherent adversarial robustness of deep spiking neural networks: Effects of discrete input encoding and non-linear activations. In *European Conference on Computer Vision*, 399–414. Springer.
- Shi, Y.; Wang, N.; and Guo, X. 2024. YOLOV: A You Only Look Once Approach to Video Object Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Su, Q.; Chou, Y.; Hu, Y.; Li, J.; Mei, S.; Zhang, Z.; and Li, G. 2023. Deep directly-trained spiking neural networks for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6555–6565.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; and Fergus, R. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Tavanaei, A.; Ghodrati, M.; Kheradpisheh, S. R.; Masquelier, T.; and Maida, A. 2019. Deep learning in spiking neural networks. *Neural networks*, 111: 47–63.
- Werbos, P. J. 2002. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10): 1550–1560.
- Wu, H.; Rowlands, S.; and Wahlström, J. 2024. A Human-in-the-Middle Attack against Object Detection Systems. *IEEE Transactions on Artificial Intelligence*.
- Wu, Y.; Deng, L.; Li, G.; Zhu, J.; and Shi, L. 2018. Spatio-temporal backpropagation for training high-performance spiking neural networks. *Frontiers in neuroscience*, 12: 331.
- Zheng, H.; Wu, Y.; Deng, L.; Hu, Y.; and Li, G. 2021. Going deeper with directly-trained larger spiking neural networks.

In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11062–11070.

Zheng, M.; Sun, L.; and Dong, J. 2025. Efficient Video Super-Resolution for Real-time Rendering with Decoupled G-buffer Guidance. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*.

Zhu, R.-J.; Zhao, Q.; Li, G.; and Eshraghian, J. K. 2023. Spikept: Generative pre-trained language model with spiking neural networks. *arXiv preprint arXiv:2302.13939*.