
Enhancing Adversarial Robustness in SNNs with Sparse Gradients

Yujia Liu¹²³ Tong Bu¹²⁴ Jianhao Ding¹²³ Zecheng Hao¹²³ Tiejun Huang¹²³ Zhaofei Yu¹²⁴

Abstract

Spiking Neural Networks (SNNs) have attracted great attention for their energy-efficient operations and biologically inspired structures, offering potential advantages over Artificial Neural Networks (ANNs) in terms of energy efficiency and interpretability. Nonetheless, similar to ANNs, the robustness of SNNs remains a challenge, especially when facing adversarial attacks. Existing techniques, whether adapted from ANNs or specifically designed for SNNs, exhibit limitations in training SNNs or defending against strong attacks. In this paper, we propose a novel approach to enhance the robustness of SNNs through gradient sparsity regularization. We observe that SNNs exhibit greater resilience to random perturbations compared to adversarial perturbations, even at larger scales. Motivated by this, we aim to narrow the gap between SNNs under adversarial and random perturbations, thereby improving their overall robustness. To achieve this, we theoretically prove that this performance gap is upper bounded by the gradient sparsity of the probability associated with the true label concerning the input image, laying the groundwork for a practical strategy to train robust SNNs by regularizing the gradient sparsity. We validate the effectiveness of our approach through extensive experiments on both image-based and event-based datasets. The results demonstrate notable improvements in the robustness of SNNs. Our work highlights the importance of gradient sparsity in SNNs and its role in enhancing robustness.

1. Introduction

Although Artificial Neural Networks (ANNs) have achieved impressive performance across various tasks (He et al., 2016; Mishra et al., 2021; Khan et al., 2023), they are often plagued by complex computations and limited interpretability (Liu et al., 2021; Lipton, 2018). In recent years, Spiking Neural Networks (SNNs) have garnered significant attention in the field of artificial intelligence due to their energy-efficient operations and biologically-inspired architectures (Maass, 1997; Zenke et al., 2021). In SNNs, neurons simulate changes in membrane potentials and transmit information through spike trains (Roy et al., 2019). These characteristics allow for a certain level of biological interpretation while avoiding the extensive and complex matrix multiplication operations inherent in ANNs. Remarkably, SNNs have achieved competitive performance with ANNs on various classification datasets (Sengupta et al., 2019; Fang et al., 2021a; Deng et al., 2022; Shao et al., 2023).

Similar to ANNs (Tanay & Griffin, 2016; Stutz et al., 2019; Liu et al., 2022), the issue of robustness poses a significant challenge for SNNs (Sharmin et al., 2019; 2020; Kundu et al., 2021). When subjected to imperceptible perturbations added to input images, SNNs can exhibit misclassifications, which are known as adversarial attacks. Developing techniques to train adversarially robust SNNs remains an ongoing problem in the research community. Some successful techniques designed for ANNs have been adapted for use with SNNs, including adversarial training (Madry et al., 2018; Ho et al., 2022; Ding et al., 2022) and certified training (Zhang et al., 2020; 2021; Liang et al., 2022). Additionally, there are SNN-specific methods proposed to improve robustness, such as temporal penalty configurations (Leontev et al., 2021), and specialized coding schemes (Sharmin et al., 2020). However, these methods either prove challenging to train on SNNs (Liang et al., 2022) or exhibit limited effectiveness against strong attacks (Sharmin et al., 2020).

In this paper, we present a novel approach to enhance the robustness of SNNs by considering the gradient sparsity with respect to the input image. We find that SNNs exhibit greater robustness to random perturbations, even at larger scales, compared to adversarial perturbations. Building upon this observation, we propose to minimize the performance gap between an SNN subjected to adversarial perturbations and

¹NERCVT, School of Computer Science, Peking University, China ²National Key Laboratory for Multimedia Information Processing, Peking University, China ³School of Computer Science, Peking University, China ⁴Institution for Artificial Intelligence, Peking University, China. Correspondence to: Zhaofei Yu <yuzf12@pku.edu.cn>.

random perturbations, thereby enhancing its overall robustness. The main contributions of our work are as follows and the code of this work is accessible at https://github.com/putshua/gradient_reg_defense.

- We analyze the robustness of SNNs and reveal that SNNs exhibit robustness against random perturbations even at significant scales, but display vulnerability to small-scale adversarial perturbations.
- We provide theoretical proof that the gap between the robustness of SNNs under these two types of perturbations is upper bounded by the sparsity of gradients of the probability associated with the true label with respect to the input image.
- We propose to incorporate gradient sparsity regularization into the loss function during training to narrow the gap, thereby boosting the robustness of SNNs.
- Extensive experiments on the image-based and event-based datasets validate the effectiveness of our method, which significantly improves the robustness of SNNs.

2. Related Work

2.1. Learning Algorithms of SNNs

The primary objective of most SNN learning algorithms is to achieve high-performance SNNs with low latency. Currently, the most effective and popular learning algorithms for SNNs are the ANN-SNN conversion (Cao et al., 2015) and supervised learning (Wu et al., 2018). The ANN-SNN conversion method aims to obtain SNN weights from pre-trained ANNs with the same network structure. By utilizing weight scaling (Li et al., 2021; Hu et al., 2023), threshold balancing (Diehl et al., 2015; Deng & Gu, 2021), quantization training techniques (Bu et al., 2022), and spike calibration (Hao et al., 2023b), well-designed ANN-SNN algorithms can achieve lossless performance compared to the original ANN (Han et al., 2020; Ho & Chang, 2021). However, the converted SNNs often require larger timesteps to achieve high performance, resulting in increased energy consumption. Moreover, they lose temporal information and struggle to process neuromorphic datasets. The supervised learning approach directly employs the backpropagation algorithm to train SNNs with fixed timesteps. Wu et al., (2018; 2019) borrowed the idea from the Back Propagation Through Time (BPTT) in RNN learning and proposed the Spatio-Temporal-Back-Propagation (STBP) algorithm. They approximate the gradient of spiking neurons using surrogate functions (Neftci et al., 2019). While supervised training significantly improves the performance of SNNs on classification tasks (Kim & Panda, 2021; Lee et al., 2020; Fang et al., 2021b; Zheng et al., 2021; Guo et al., 2022; Yao et al., 2022; Duan et al., 2022; Mostafa, 2017; Bohte et al., 2000; Zhang & Li, 2020; Zhang et al., 2022; Xu et al.,

2022b; Zhu et al., 2022; 2023), SNNs still fall behind ANNs in terms of generalization and flexibility. Challenges such as gradient explosion and vanishing persist in SNNs.

2.2. Defense Methods of SNNs

Methods for improving the robustness of SNNs can be broadly categorized into two classes. The first class draws inspiration from ANNs. A typical representative is adversarial training, which augments the training set with adversarial examples generated by attacks (Madry et al., 2018; Tramèr et al., 2018; Wong et al., 2020). This approach has been shown to effectively defend against attacks that are used in the training phase. Another method is certified training, which utilizes certified defense methods to train a network (Wong & Kolter, 2018; Xu et al., 2020). Certified training has demonstrated promising improvements in the robustness of ANNs (Zhang et al., 2020), but its application to SNNs remains challenging (Liang et al., 2022). The second category consists of SNN-specific techniques designed to enhance robustness. On one hand, the choice of encoding the continuous intensity of an image into 0-1 spikes can impact the robustness of SNNs. Recent studies have highlighted the Poisson encoder as a more robust option (Sharmin et al., 2020; Kim et al., 2022). However, the Poisson encoder generally yields worse accuracy on clean images than the direct encoding, and the robustness improvement caused by the Poisson encoder varies with the number of timesteps used. On the other hand, researchers have recognized the unique temporal dimension of SNNs and developed strategies related to temporal aspects to improve robustness (Nomura et al., 2022). Hao et al. (2023a) further pointed out the significance of utilizing the rate and temporal information comprehensively to enhance the reliability of SNNs. Apart from the studies on static datasets, there are works that attempt to perform adversarial attacks and defenses on the Dynamic Vision Sensors (DVS) dataset (Marchisio et al., 2021b). In this paper, we mainly focus on the direct encoding of input images. We propose a gradient sparsity regularization strategy to improve SNNs’ robustness with theoretical guarantees. Moreover, this strategy can be combined with adversarial training to further boost the robustness of SNNs.

3. Preliminary

3.1. Neuron Dynamics in SNNs

Similar to previous works (Wu et al., 2018; Rathi & Roy, 2021), we consider the commonly used Leaky Integrate-and-Fire (LIF) neuron model due to its efficiency and simplicity, the dynamic of which can be formulated as follows:

$$u_i^l[t] = \tau u_i^l[t-1](1 - s_i^l[t-1]) + \sum_j w_{i,j}^{l-1} s_j^{l-1}[t], \quad (1)$$

$$s_i^l[t] = H(u_i^l[t] - \theta). \quad (2)$$

Equation (1) describes the membrane potential of the i -th neuron in layer l , which receives the synaptic current from the j -th neuron in layer $l - 1$. Here τ represents the membrane time constant and t denotes the discrete time step ranging from 1 to T . The variable $u_i^l[t]$ represents the membrane potential of the i -th neuron in layer l at the time step t . $w_{i,j}$ denotes the synaptic weight between the two neurons, and $s_j^{l-1}[t]$ represents the binary output spike of neuron j in layer $l - 1$. For simplicity, the resting potential is assumed to be zero so that the membrane potential will be reset to zero after firing. Equation (2) defines the neuron fire function. At each time step t , a spike will be emitted when the membrane potential $u_i^l[t]$ surpasses a specific threshold θ . The function $H(\cdot)$ denotes the Heaviside step function, which equals 0 for negative input and 1 for others.

3.2. Adversarial Attacks for SNNs

Preliminary explorations have revealed that SNNs are also susceptible to adversarial attacks (Sharmin et al., 2020; Kundu et al., 2021; Liang et al., 2021; Marchisio et al., 2021a). Well-established techniques such as the Fast Gradient Sign Method (FGSM) and Projected Gradient Descent (PGD) can generate strong adversaries that threaten SNNs.

FGSM (Goodfellow et al., 2015) is a straightforward non-iterative attack method, expressed as follows:

$$\hat{x} = x + \epsilon \text{sign}(\nabla_x \mathcal{L}(f(x), y)), \quad (3)$$

where x and \hat{x} represent the original image and the adversarial example respectively, ϵ denotes the perturbation bound, \mathcal{L} refers to the loss function, $f(\cdot)$ represents the neural network function, and y denotes the label data.

PGD (Madry et al., 2018) is an iterative extension of FGSM, which can be described as follows:

$$\hat{x}^k = \Pi_\epsilon \{x^{k-1} + \alpha \text{sign}(\nabla_x \mathcal{L}(f(x^{k-1}), y))\}, \quad (4)$$

where k is the current iteration step and α is the step size. The operator Π_ϵ projects the adversarial examples onto the space of the ϵ neighborhood in the ℓ_∞ norm around x .

4. Methodology

In this section, we first compare the vulnerability of SNNs to random perturbations versus adversarial perturbations. We highlight that SNNs exhibit significant robustness against random perturbations but are more susceptible to adversarial perturbations. Then we quantify the disparity between adversarial vulnerability and random vulnerability, proving that it is upper bounded by the gradient sparsity of the probability related to the true label concerning the input image. Based on this, we propose a novel approach to enhance the robustness of SNNs by introducing Sparsity Regularization (SR) of gradients in the training phase and incorporating this regularization into the learning rule of SNNs.

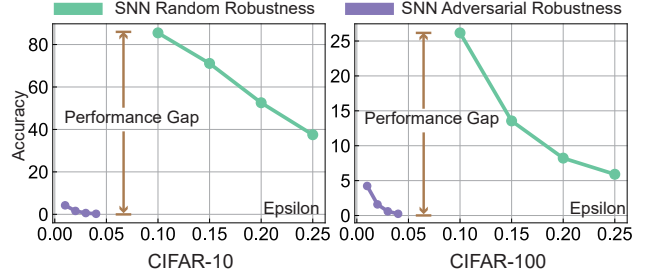


Figure 1: Comparison of the random vulnerability and adversarial vulnerability of SNNs on CIFAR-10 and CIFAR-100.

4.1. Compare Random and Adversarial Vulnerability

The objective of adversarial attacks is to deliberately alter the probability vector $f(x)$ in order to change the classification result. Since the classification result is determined by the magnitude-based ordering of components in $f(x)$, attackers aim to substantially decrease the value of the component corresponding to the true label of x . We suppose x with label y belonging to the y -th class, so that the value of $f_y(x)$ has a substantial impact on the classification result. Therefore, the stability of $f_y(x)$ becomes crucial for the robustness of SNNs, particularly in terms of the value of $|f_y(\hat{x}) - f_y(x)|$ when subjected to small perturbation on x .

To initiate our analysis, we define the random and adversarial vulnerability of an SNN, denoted as f , at a specific point x under an ℓ_p attack of size ϵ .

Definition 4.1. (Random Vulnerability) The random vulnerability of f at point x to an ℓ_p attack of size ϵ is defined as the expected value of $(f_y(x + \epsilon \cdot \delta) - f_y(x))^2$, where δ follows a uniform distribution within the unit ℓ_p ball, and y represents the class of x belonging to. Mathematically, it can be expressed as:

$$\rho_{\text{rand}}(f, x, \epsilon, \ell_p) = \mathbb{E}_{\delta \sim U\{\|\delta\|_p \leq 1\}} (f_y(x + \epsilon \cdot \delta) - f_y(x))^2. \quad (5)$$

Definition 4.2. (Adversarial Vulnerability) The adversarial vulnerability of f at point x to an ℓ_p attack of size ϵ is defined as the supremum of $(f_y(x + \epsilon \cdot \delta) - f_y(x))^2$, where δ follows a uniform distribution within the unit ℓ_p ball, and y represents the class of x belonging to. Mathematically, it can be expressed as:

$$\rho_{\text{adv}}(f, x, \epsilon, \ell_p) = \sup_{\delta \sim U\{\|\delta\|_p \leq 1\}} (f_y(x + \epsilon \cdot \delta) - f_y(x))^2. \quad (6)$$

To gain a deeper understanding of the disparity in vulnerability between random perturbations and adversarial perturbations in SNNs, we conduct a small-scale experiment with a primary focus on ℓ_∞ attacks. The results of the experiment are depicted in Figure 1. We specifically prioritized ℓ_∞ attacks over other ℓ_p attacks due to the widespread

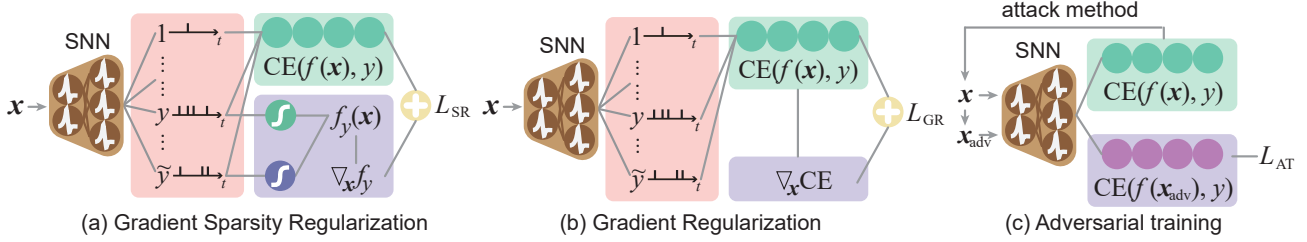


Figure 2: Illustration of (a) the proposed SR strategy, (b) gradient regularization and (c) adversarial training.

utilization of ℓ_∞ constraints in various attack methods. Adversarial examples generated under ℓ_∞ attacks tend to be more destructive compared to those generated under ℓ_0 and ℓ_2 attacks (Madry et al., 2018). Therefore, focusing on ℓ_∞ attacks allows us to assess the robustness of the SNN against the most severe random and adversarial perturbations.

In our experiment, we evaluate the performance of a well-trained SNN f with the VGG-11 architecture. Our primary objective is to examine the impact of both random and adversarial perturbations on the SNN’s classification results. For a given image x , we added random perturbations uniformly drawn from a hyper-cube $\{\delta_{\text{rand}} : \|\delta_{\text{rand}}\|_\infty \leq \epsilon\}$ to the original image, resulting in the perturbed image \hat{x}_{rand} . Additionally, we employed a ϵ -sized FGSM attack to generate an adversarial example \hat{x}_{adv} from x . Subsequently, we individually input both \hat{x}_{rand} and \hat{x}_{adv} into the network f to observe any changes in the classification results. We evaluate the classification accuracy of the perturbed images on the test set of CIFAR-10 and CIFAR-100 under both random and adversarial perturbations, respectively. The results, as depicted in Figure 1, yield the following key findings:

Observation 1. *SNNs exhibit robustness against random perturbations even when the perturbation scale is significant, but display vulnerability to small-scale adversarial perturbations.*

4.2. Quantify the Gap between Random Vulnerability and Adversarial Vulnerability

Observation 1 indicates that SNNs exhibit notable robustness to random perturbations in comparison to adversarial perturbations. To enhance the adversarial robustness of SNNs, a natural approach is to minimize the disparity between adversarial vulnerability and random vulnerability.

To measure the disparity in vulnerability between the SNN f under adversarial perturbations and random noise, we employ the ratio of $\rho_{\text{adv}}(f, x, \epsilon, \ell_\infty)$ and $\rho_{\text{rand}}(f, x, \epsilon, \ell_\infty)$. We make the assumption that $\rho_{\text{rand}}(f, x, \epsilon, \ell_\infty) \neq 0$, indicating that f is not a constant function. This assumption aligns with the practical reality of SNNs in real-world applications. Optimizing the ratio of $\rho_{\text{adv}}(f, x, \epsilon, \ell_\infty)$ and $\rho_{\text{rand}}(f, x, \epsilon, \ell_\infty)$ directly is a challenging task. However,

we are fortunate to present a mathematical proof that establishes an upper bound for this ratio based on the sparsity of $\nabla_x f_y$. Specifically, we have the following theorem.

Theorem 4.3. *Suppose f is a differentiable SNN by surrogate gradients, and ϵ is the magnitude of an attack, assumed to be small enough. Given an input image x with corresponding label y , the ratio of $\rho_{\text{adv}}(f, x, \epsilon, \ell_\infty)$ and $\rho_{\text{rand}}(f, x, \epsilon, \ell_\infty)$ is upper bounded by the sparsity of $\nabla_x f_y$:*

$$3 \leq \frac{\rho_{\text{adv}}(f, x, \epsilon, \ell_\infty)}{\rho_{\text{rand}}(f, x, \epsilon, \ell_\infty)} \leq 3 \|\nabla_x f_y(x)\|_0. \quad (7)$$

The proof is provided in Appendix A. This theorem illustrates that the disparity between the adversarial vulnerability and random vulnerability is upper bounded by the sparsity of $\nabla_x f_y$. It provides valuable insights into the correlation between gradient sparsity and the disparity in robustness exhibited by SNNs when subjected to different perturbations with ℓ_∞ attacks. According to Theorem 4.3, we can infer that a sparser gradient contributes to closing the robustness gap between SNN f under worst-case scenarios and its robustness under random perturbations.

From an intuitive perspective, minimizing the ℓ_0 norm of $\nabla_x f_y(x)$ serves to bring x closer to a local minimum point. In an ideal scenario, this would entail trapping x within a local minimum, effectively rendering attackers unable to generate adversarial examples through gradient-based methods. By introducing the sparsity constraint for each x in the training set, we encourage learning an SNN f where input images tend to remain close to extreme points or trapped in local minimums. This makes it challenging to perturb $f_y(x)$ with small perturbations, thereby enhancing the robustness of the SNN f .

4.3. Loss Function with Sparsity Regularization

To promote sparsity of the gradients, a straightforward approach is to incorporate the $\|\nabla_x f_y(x)\|_0$ term into the training loss, where $f_y(x)$ is the probability assigned by f to x belonging to the true label y .

Consider an SNN f with a final layer denoted as L . The total number of neurons in layer L is denoted by N , and the

time-step t ranges from 1 to T . For a given input image \mathbf{x} , the output vector of the L^{th} layer depends on the collective outputs across all time-steps:

$$f^L(\mathbf{x}) = \left(\sum_{t=1}^T s_1^L(t), \dots, \sum_{t=1}^T s_N^L(t) \right)^T. \quad (8)$$

While regularizing $\|\nabla_{\mathbf{x}} f_y^L(\mathbf{x})\|_0$ is a straightforward way to keep the stability of $f_y^L(\mathbf{x})$, it may be insufficient for multi-classification tasks. This is because the classification result is influenced not only by the value of f_y^L , but also by the magnitude of f_y^L in comparison to the other components of f_y^L . To account for such relationships while maintaining computational efficiency, we utilize two spike streams at the last layer of f to calculate $f_y(\mathbf{x})$, as illustrated in Figure 2 (a). This is expressed as:

$$\begin{aligned} f_y &= \frac{e^{\sum_{t=1}^T s_{\tilde{y}}^L(t)}}}{e^{\sum_{t=1}^T s_y^L(t)} + e^{\sum_{t=1}^T s_{\tilde{y}}^L(t)}}, \\ \tilde{f}_{\tilde{y}} &= \frac{e^{\sum_{t=1}^T s_{\tilde{y}}^L(t)}}}{e^{\sum_{t=1}^T s_y^L(t)} + e^{\sum_{t=1}^T s_{\tilde{y}}^L(t)}}. \end{aligned} \quad (9)$$

where \tilde{y} represents the index of the maximum component in $\{f_i^L(\mathbf{x}) : i \neq y\}$.

On one hand, the transformation in Equation (9) introduces only one additional component in f^L while preserving the classification results derived from f^L , as expressed by

$$\arg \max_{i=1, \dots, N} f_i^L(\mathbf{x}) = \arg \max_{i=y, \tilde{y}} f_i(\mathbf{x}). \quad (10)$$

On the other hand, since $f_y + \tilde{f}_{\tilde{y}} = 1$, so if f_y is stable, $\tilde{f}_{\tilde{y}}$ is also stable. Thus, it is sufficient to regularize $\|\nabla_{\mathbf{x}} f_y(\mathbf{x})\|_0$ to enhance the adversarial robustness of the SNN f .

Therefore, the training loss can be written as:

$$\mathcal{L}(\mathbf{x}, y) = \text{CE}(f^L(\mathbf{x}), y) + \lambda \|\nabla_{\mathbf{x}} f_y(\mathbf{x})\|_0, \quad (11)$$

where $\text{CE}(\cdot)$ is the cross-entropy loss, $f^L(\mathbf{x})$ is the output of the last layer, and λ denotes the coefficient parameter controlling the strength of the sparsity regularization.

Here we give the formulation of $\nabla_{\mathbf{x}} f_y(\mathbf{x})$ and the detailed derivation is provided in Appendix B. Given an input image \mathbf{x} belonging to the class y , $\nabla_{\mathbf{x}} f_y(\mathbf{x})$ can be formulated as:

$$\nabla_{\mathbf{x}} f_y(\mathbf{x}) = \sum_{i=y, \tilde{y}} \left(\frac{\partial f_y}{\partial f_i^L} \left(\sum_{t=1}^T \sum_{i=1}^t \nabla_{\mathbf{x}[\tilde{y}]} s_i^L[t] \right) \right). \quad (12)$$

It is worth noting that the optimization problem involving the ℓ_0 norm is known to be NP-hard (Natarajan, 1995). To circumvent this challenge, we employ the ℓ_1 norm as a substitute for the ℓ_0 norm because the ℓ_1 norm serves as a

Algorithm 1 Training Algorithm

Input: Spiking neural network $f(\mathbf{x}, w)$ with parameter w
 Learning rate η ; Step size h of finite differences; Balance weight λ

Output: Regularize trained parameter w

```

1: for epoch=0 to n do
2:   Sample minibatch  $\{(\mathbf{x}^i, y^i)\}_{i=1, \dots, m}$  from Dataset
3:   for  $i = 0$  to  $m$  do
4:      $f^L = f(\mathbf{x}^i, w)$ 
5:      $\tilde{y}^i = \arg \max_{j \neq y} f_j^L$ 
6:      $f_{y^i} = e^{f_{y^i}^L} / (e^{f_{y^i}^L} + e^{f_{\tilde{y}^i}^L})$ 
7:      $\mathbf{d}^i = \text{sign}(\nabla_{\mathbf{x}} f_{y^i}) \leftarrow$  the difference direction
8:      $\hat{\mathbf{x}}^i = \mathbf{x}^i + h \mathbf{d}^i$ 
9:      $\mathcal{L}(\mathbf{x}^i, y^i, w) = \text{CE}(f^L, y^i) + \frac{\lambda}{h} |f_{y^i}(\hat{\mathbf{x}}^i) - f_{y^i}(\mathbf{x}^i)|$ 
10:     $w \leftarrow w - \eta \nabla_w \mathcal{L}(\mathbf{x}^i, y^i, w)$ 
11:   end for
12: end for
    
```

convex approximation to the ℓ_0 norm (Ramirez et al., 2013). However, the computational burden associated with calculating the back-propagation of $\|\nabla_{\mathbf{x}} f_y(\mathbf{x})\|_1$ is significant. Meanwhile, we find that SNNs are not trainable with such a double backpropagation approach. Therefore, we adopt a finite difference approximation for this term (Finlay & Oberman, 2021). In our case, we approximate the gradient regularization term using the following finite differences:

Proposition 4.4. *Let \mathbf{d} denote the signed input gradient direction: $\mathbf{d} = \text{sign}(\nabla_{\mathbf{x}} f_y(\mathbf{x}))$, and h be the finite difference step size. Then, the ℓ_1 gradient norm can be approximated as:*

$$\|\nabla_{\mathbf{x}} f_y(\mathbf{x})\|_1 \approx \left| \frac{f_y(\mathbf{x} + h \cdot \mathbf{d}) - f_y(\mathbf{x})}{h} \right|. \quad (13)$$

The proof is provided in the Appendix C. Finally, the training loss (Equation (11)) is rewritten as:

$$\mathcal{L}(\mathbf{x}, y) = \text{CE}(f^L(\mathbf{x}), y) + \lambda \left| \frac{f_y(\mathbf{x} + h \cdot \mathbf{d}) - f_y(\mathbf{x})}{h} \right|. \quad (14)$$

The overall training algorithm is presented as Algorithm 1.

4.4. Differences with Related Works

We compare the proposed Sparsity Regularization (SR) strategy with Gradient Regularization (GR) as proposed by (Finlay & Oberman, 2021) and classic adversarial training (Madry et al., 2018) in Figure 2. While GR relies on the selection of the multi-class calibrated loss function (e.g., cross-entropy or logistic loss), and adversarial training is associated with the attack method used in generating adversarial examples, the proposed SR strategy is both loss-independent and attack-independent.

Table 1: Comparison with the SOTA methods on classification accuracy (%) under attacks.

Dataset	Arch.	Defense	Clean	White Box Attack				Black Box Attack			
				PGD10	PGD30	PGD50	APGD10	PGD10	PGD30	PGD50	APGD10
CIFAR-10	VGG-11	RAT	90.44	11.53	7.08	6.41	3.26	43.29	40.17	40.16	47.50
		AT	89.97	18.18	14.79	14.63	10.36	44.02	43.38	43.40	52.90
		SR*	85.91	30.54	28.06	27.66	21.91	51.21	50.85	51.06	59.87
CIFAR-10	WRN-16	RAT	92.70	10.52	5.14	4.33	2.19	38.35	31.04	30.40	36.42
		AT	90.97	17.88	14.89	14.62	9.13	43.99	42.09	41.55	52.10
		SR*	85.63	39.18	37.04	36.74	29.03	50.84	50.23	49.97	57.93
CIFAR-100	WRN-16	RAT	69.10	5.72	3.58	3.26	2.08	22.61	18.77	18.26	25.23
		AT	67.37	10.07	8.12	7.86	4.88	25.17	23.76	23.50	35.96
		SR*	60.37	19.76	18.39	18.11	13.32	28.38	28.01	27.94	36.69

SR Strategy vs. GR Strategy. The main distinction between SR and GR strategies lies in that the regularization term in SR is exclusively tied to the *model itself*, rather than being dependent on the multi-class calibrated loss used during the training phase, as shown in Figure 2 (b). To provide further clarification, we can express the training loss of the GR strategy as follows (defending against attacks):

$$\ell(f^L(x), y) + \lambda \|\nabla_x \ell(f^L(x), y)\|_1^2 \quad (15)$$

where $\ell(f^L(x), y)$ is the multi-class calibrated loss function for the classification task (such as cross-entropy loss). Note that the regularization term in GR varies depending on the specific choices of multi-class calibrated loss. In contrast, our proposed SR method introduces a regularization term that remains independent of the choice of loss function. Moreover, our method can be combined with adversarial training to enhance the performance of adversarial training.

SR Strategy vs. Adversarial Training. The key distinction between these two strategies lies in their approach. Adversarial training involves generating adversarial examples using specific attack methods, while the SR strategy is independent of adversarial examples. Figure 2 (c) illustrates that adversarial training aims to minimize the multi-class calibrated loss for adversarial images generated by attacks such as FGSM or PGD. But the main idea of SR revolves around regularizing the sparsity of $\nabla_x f_y(\mathbf{x})$. Although we employ the finite difference method (Equation (14)), to make the regularization of $\|\nabla_x f_y(\mathbf{x})\|_1$ computationally feasible, it is essential to recognize that the sample $\mathbf{x} + h\mathbf{d}$ fundamentally differs from an adversarial example in both numerical value and meaning. On one hand, $\mathbf{x} + h\mathbf{d}$ is employed to calculate the difference quotient instead of directly calculating the multi-class calibrated loss. On the other hand, since h serves as the step size for the approximation, it is advisable to use a small value to obtain a more accurate estimation of the ℓ_1 norm. This contrasts with the requirement for a large h when generating adversarial examples.

5. Experiment

In this section, we evaluate the performance of the proposed SR strategy on image classification tasks using the CIFAR-10, CIFAR-100 and CIFAR10-DVS datasets. We adopt the experiment setting used in the previous work (Ding et al., 2022). Specifically, we use the VGG-11 architecture (Simonyan & Zisserman, 2014), WideResNet with a depth of 16 and width of 4 (WRN-16) (Zagoruyko & Komodakis, 2016). The timestep for the SNNs is set to 8. Throughout the paper, we use the IF neuron with a hard-reset mechanism as the spiking neuron. Further details regarding the training settings can be found in the Appendix D.

To generate adversarial examples, we employ different attacks, including FGSM (Goodfellow et al., 2015), PGD (Madry et al., 2018), and AutoPGD (Croce & Hein, 2020), with a fixed attack strength of 8/255. For iterative attacks, the number of iterations is indicated in the attack name (e.g. PGD10). Since the choice of gradient approximation methods (Bu et al., 2023) and surrogate functions can affect the attack success rate (Xu et al., 2022a), we consider an ensemble attack for SNNs (Özdenizci & Legenstein, 2023). We utilize a diverse set of surrogate gradients and consider both STBP-based (Esser et al., 2016) and RGA-based (Bu et al., 2023) attacks. For each test sample, we conduct multiple attacks using all possible combinations of gradient approximation methods and surrogate functions, and report the strongest attack. In other words, we consider an ensemble attack to be successful for a test sample as long as the model is fooled with any of the attacks from the ensemble. Robustness is evaluated in two scenarios: the white-box scenario, where attackers have knowledge of the target model, and the black-box scenario, where the target model is unknown to attackers. More detailed evaluation settings can be found in the Appendix E. Moreover, for the ensemble attack to be meaningful and reveal any impact of gradient obfuscation, we run extensive experiment with varying widths of surrogate functions in Appendix G.

5.1. Compare with the State-of-the-art

We validate the effectiveness of our method by comparing it with the current state-of-the-art approaches, including Regularized Adversarial Training (RAT) (Ding et al., 2022) and Adversarial Training (AT) (Kundu et al., 2021). We use SR* to denote our sparsity regularization strategy with adversarial training. For RAT-SNN, we replicate the model following the settings outlined in the paper (Ding et al., 2022). As for all SNNs trained with robustness training strategies, we adopt a PGD5 attack with $\epsilon = 2/255$.

Table 1 reports the classification accuracy of the compared methods under ensemble attacks. Columns 5-8 highlight the substantial enhancement in SNN robustness achieved through our strategy in the white box scenario. Our proposed method consistently outperforms other State-Of-The-Art (SOTA) methods across all datasets and architectures. For instance, when subjected to 10-steps PGD attacks, VGG-11 trained with our strategy elevates classification accuracy from 11.53% (RAT) to an impressive 30.54% on CIFAR-10. Similarly, WRN-16 trained with SR* exhibits a remarkable 15% boost in classification accuracy against PGD50 on CIFAR-100. In comparison to the AT strategy, our method demonstrates a noteworthy enhancement in adversarial robustness, with a 10-20 percentage point improvement on both datasets under all attacks.

In contrast to white box attacks, all strategies exhibit better adversarial robustness against black box attacks (columns 9-12 in Table 1). When considering the CIFAR-10 dataset, models trained with any strategy achieve a classification accuracy of over 30% when subjected to PGD50. However, models trained with the SR* strategy consistently outperform other strategies in all scenarios. There exists a gap of 10%-20% in performance between models trained with RAT/AT and those trained with SR* on CIFAR-10, and a 5%-10% gap on CIFAR-100. These results demonstrate the superiority of our approach over SOTA methods.

5.2. Experiments on Dynamic Vision Sensor Data

Since SNNs are suitable for application on neuromorphic data, we evaluate the effectiveness of the gradient sparsity regularization on CIFAR10-DVS dataset. Here, we use the preprocessed neuromorphic data and each data point contains ten frame-based data. The data batches are then fed into a 10-timestep spiking neural network for training and inference and we directly generate the adversarial noise on the frame-based data. Similar to previous experiments, we choose VGG-11 architecture and compare models trained from defense approaches including vanilla model, adversarial trained model and SR* strategy trained model. Here the regularization coefficient parameter is set to $\lambda = 0.002$.

Table 2 demonstrate the performance comparison under ad-

Table 2: Experiments on CIFAR10-DVS.

Defense	Clean	White Box Attack		Black Box Attack	
		FGSM	PGD50	FGSM	PGD50
Vanilla	78.80	22.20	4.40	34.00	20.70
AT	76.80	60.00	51.60	71.50	65.50
SR	77.40	37.30	27.90	44.60	37.70
SR*	75.60	64.60	61.20	72.60	68.90

Table 3: Ablation study of the sparsity regularization.

SR	AT	Clean	FGSM	RFGSM	PGD30	PGD50	APGD10
		CIFAR-10			WRN-16		
×	×	93.89	5.23	3.43	0.00	0.00	0.00
×	✓	90.97	33.49	58.19	14.89	14.62	9.13
✓	×	86.57	34.79	55.96	12.27	11.70	8.25
✓	✓	85.63	48.47	64.65	37.04	36.74	29.03
		CIFAR-100			WRN-16		
×	×	74.59	3.51	1.37	0.00	0.00	0.00
×	✓	67.37	19.07	33.19	8.12	7.68	4.88
✓	×	67.67	11.15	18.18	0.87	0.84	0.47
✓	✓	60.37	25.76	36.93	18.39	18.11	13.32

versarial attack with $\epsilon = 0.031$. As can be seen from the table, the vanilla SNN can be easily fooled by the PGD50 attack and the robust performance is only 4.4% under white-box PGD50 attack. However, with the application of adversarial training and gradient sparsity regularization, the robust performance shows a significant improvement. The adversarially trained model’s robustness increases to 51.6% and further rises to 61.2% under PGD50 white-box attack. The SR* strategy model also exhibits a substantial improvement in adversarial robustness under black-box attacks, elevating the classification accuracy from 20.7% to 68.9% when subjected to the PGD50. These successful defenses of the sparse gradient method on the Dynamic Vision Sensor dataset proves the effectiveness and possibility of application of our method on neuromorphic datasets.

5.3. Ablation Study of the Sparsity Regularization

In the ablation study, we compare the robustness performance of SNNs using different training strategies: vanilla SNN, SR-SNN, AT-SNN, and SR*-SNN. The results on CIFAR-10 and CIFAR-100 are presented in Table 3, and the key findings are summarized as follows.

Firstly, it is crucial to note that vanilla SNNs exhibit poor adversarial robustness, with their classification accuracy dropping to a mere 5% when subjected to the FGSM attack. However, the SR strategy significantly enhances this

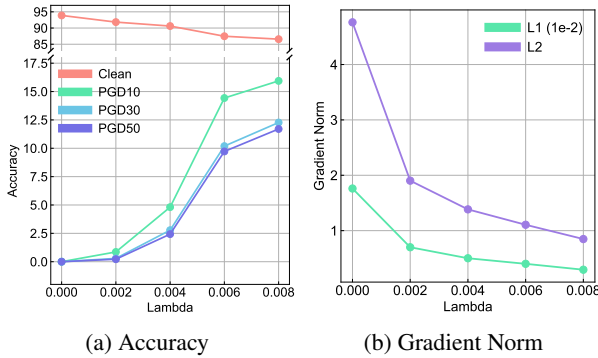


Figure 3: The influence of the coefficient parameter λ on classification accuracy and gradient sparsity. (a): Fluctuations in clean accuracy and adversarial accuracy under PGD attacks across different values of λ . (b): The ℓ_1 and ℓ_2 norms of the gradient with varying λ .

performance, achieving a classification accuracy of 34.79% on the CIFAR-10 dataset and 11.15% on the CIFAR-100 dataset. Furthermore, combining the SR strategy with the AT strategy further boosts the robustness of SNNs, particularly against strong attacks like PGD50 and APGD10, resulting in a notable 10%-30% improvement.

Additionally, it is observed that the classification accuracy of robust SNNs on clean images may typically be slightly lower than that of baseline models. This phenomenon is consistent across all robustness training strategies. For example, WRN-16 models trained using any strategy exhibit a classification accuracy of less than 70% on clean images in CIFAR-100. Striking a balance between adversarial robustness and classification accuracy on clean images remains an open challenge in the field, warranting further exploration.

5.4. Search for the Optimal Coefficient Parameter

We conduct an extensive exploration to determine the optimal coefficient parameter λ , trying to strike a balance between robustness on adversarial images and classification accuracy on clean images. The investigation specifically targets the CIFAR-10 dataset and the SNN model employs the WRN-16 architecture.

As described in Figure 3 (a), we test the impact of λ varying within the range of 0.000 to 0.008. Notably, increasing the value of λ led to a decrease in classification accuracy on clean images but a significant improvement in adversarial robustness. To be specific, when using a coefficient parameter of $\lambda = 0.008$, the classification accuracy under PGD10 attack increases from zero to 16%, while maintaining over 85% accuracy on clean images.

Figure 3 (b) provides a visual representation of the effect

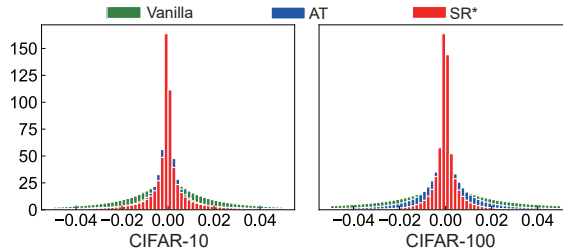


Figure 4: The normalized distribution of $\nabla_x f_y(x)$.

of λ on gradient sparsity after training. We computed the average ℓ_1 and ℓ_2 norm of $\nabla_x f_y$ over the test dataset using models trained with different λ . Values of both the ℓ_1 norm and ℓ_2 decrease significantly as the coefficient parameter increases, indicating the correctness of the approximation method introduced in Proposition 4.4 and the effectiveness of the gradient sparsity regularization.

Based on these findings, we select $\lambda = 0.008$ to train the SR-WRN-16 on the CIFAR-10 dataset to strike a balance between clean accuracy and adversarial robustness. It is worth noting that the optimal choice of λ may vary for different datasets. For additional insights into the relationship between λ , clean accuracy, and adversarial robustness on the CIFAR-100 dataset, please refer to the line chart presented in Appendix F.

5.5. Visualization of Gradient Sparsity

To validate the effectiveness of the proposed approximation method (Proposition 4.4), we compute the gradient $\nabla_x f_y(x)$ at x in three cases: f is a vanilla SNN, f is an SNN with adversarial training (AT), and f is an SNN trained with the proposed gradient sparsity regularization and adversarial training (SR*). Figure 4 illustrates the overall distribution of components in the gradient across all test samples in the CIFAR-10 and CIFAR-100 datasets, respectively.

The results clearly show that the distribution of gradient components' values for SR*-SNNs is more concentrated around zero compared to that of vanilla SNNs and AT-SNNs. This indicates that SR*-SNNs exhibit sparser gradients with respect to the input image, demonstrating the effectiveness of the finite difference method proposed in Proposition 4.4 in constraining gradient sparsity. Meanwhile, these findings suggest a correlation between the sparsity of gradients and the robustness of SNNs to some extent: sparser gradients contribute to the enhancement of SNN robustness.

To further substantiate the claim that SR-SNNs possess sparser gradients compared to vanilla SNNs, we present heatmaps of $\nabla_x f_y$ for several examples from CIFAR-10 (Tsipras et al., 2019). In Figure 5, the first row displays some original images selected from CIFAR-10, while the

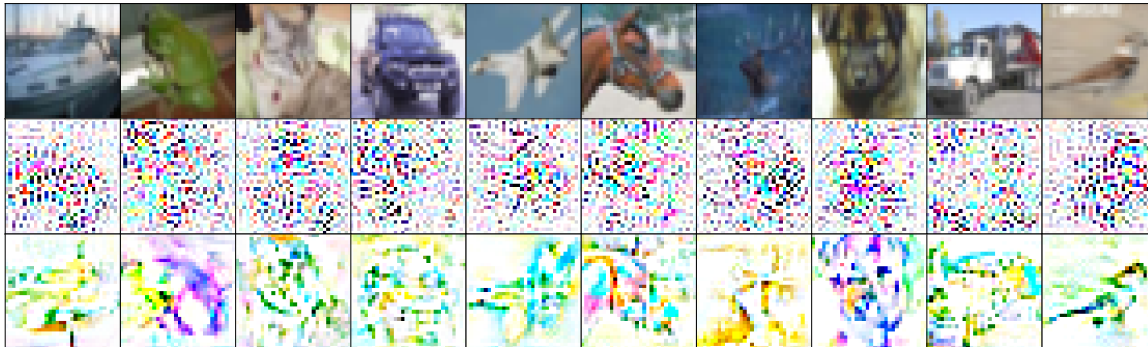


Figure 5: Heatmaps of $\nabla_x f_y$ where f is a villain SNN (top) or an SR-SNN (down).

second and third rows show the corresponding heatmaps of $\nabla_x f_y$ for vanilla SNN and SR-SNN, respectively.

Notably, the points on the heatmap of the vanilla SNN are densely distributed, while the points on the heatmap of the SR-SNN are more sparsely arranged. Moreover, the heatmap of gradients of the vanilla SNN appears cluttered to reflect any information about the image. However, the heatmap of the gradient of the SR-SNN shows some clear texture information of the image, which is beneficial to the interpretability of SNN. Therefore, we infer that the gradient sparsity regularization can not only improve the robustness of SNNs, but also provide some interpretability for SNNs.

5.6. Impact of SR strategy to the Robustness Under Random Attacks

We also investigate the impact of the SR strategy on the robustness of SNNs under random attacks. Experiments are conducted on CIFAR datasets using the WRN-16 architecture as the baseline. Random perturbations are uniformly drawn from a hyper-cube $\{\delta_{\text{rand}} : \|\delta_{\text{rand}}\|_{\infty} \leq \epsilon\}$, with $\epsilon = 0.1$. The classification accuracy of the vanilla WRN-16, SR-trained WRN-16, and SR*-trained (PGD5+SR) WRN-16 under random attacks is presented in Table 4.

According to the table, both the SR and SR* strategies significantly enhance the robustness of SNNs against random attacks. For example, the single SR strategy improves the classification accuracy by nearly 20% on the CIFAR-10 dataset. When combined with adversarial training, the SR* strategy still increases the random robustness, achieving 81% classification accuracy. This indicates that the SR strategy does not compromise robustness under random attacks while narrowing the gap between adversarial and random vulnerabilities.

In addition to the experiments reported in the main manuscript, we also analyze the computational cost of the SR strategy in Appendix H. We find that the training time for the SR strategy is less than that of PGD5 adversarial

Table 4: Classification accuracy (%) of models trained with different methods under random attacks on CIFAR datasets.

	Vanilla	SR	SR*
CIFAR-10	67.467	86.327	81.885
CIFAR-100	26.270	49.900	48.678

training, indicating that the SR strategy is efficient. More detailed explanations can be found in the appendix.

6. Conclusion and Limitation

Conclusion. This paper introduces a new perspective on SNN robustness by incorporating the concept of gradient sparsity. We theoretically prove that the ratio of adversarial vulnerability to random vulnerability in SNNs is upper bounded by the sparsity of the true label probability with respect to the input image. Moreover, we propose the SR training strategy to train robust SNNs against adversarial attacks. Experimental results confirm the consistency between the theoretical analysis and practical tests. This insight is expected to inspire ongoing research focused on enhancing SNN robustness by reducing gradient sparsity. Furthermore, it may also spark interest in investigating the robustness of event-driven SNNs, which naturally exhibit strong spike sparsity.

Limitation. The limitation of this work is that the improvement in adversarial robustness achieved by the SR strategy comes at the cost of a notable accuracy loss on clean images. In future work, we aim to strike a better balance between classification accuracy and adversarial robustness. For instance, we may employ the simulated annealing algorithm or structure learning (Bellec et al., 2018; Shen et al., 2023) in the SR strategy. Additionally, considering that biological vision has strong robustness (Dapello et al., 2020; Yu et al., 2024), proposing new SNN models by drawing inspiration from the mechanisms of biological vision is an important direction for the future.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (62176003, 62088102) and the Beijing Nova Program (20230484362). We also acknowledge Lei Wu for valuable discussions.

Impact Statement

Our research makes contributions to the field of SNN by addressing the importance of gradient sparsity. This work provides both advanced theoretical understanding and practical solutions for robust SNN training, with a deliberate focus on ensuring no discernible negative societal consequences.

References

- Bellec, G., Kappel, D., Maass, W., and Legenstein, R. Deep rewiring: Training very sparse deep networks. In *International Conference on Learning Representations*, pp. 1–24, 2018.
- Bohte, S. M., Kok, J. N., and La Poutré, J. A. Spikeprop: Backpropagation for networks of spiking neurons. In *European Symposium on Artificial Neural Networks*, pp. 419–424, 2000.
- Bu, T., Fang, W., Ding, J., Dai, P., Yu, Z., and Huang, T. Optimal ANN-SNN conversion for high-accuracy and ultra-low-latency spiking neural networks. In *International Conference on Learning Representations*, pp. 1–19, 2022.
- Bu, T., Ding, J., Hao, Z., and Yu, Z. Rate gradient approximation attack threatens deep spiking neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7896–7906, 2023.
- Cao, Y., Chen, Y., and Khosla, D. Spiking deep convolutional neural networks for energy-efficient object recognition. *International Journal of Computer Vision*, 113(1): 54–66, 2015.
- Croce, F. and Hein, M. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International Conference on Machine Learning*, pp. 2206–2216, 2020.
- Dapello, J., Marques, T., Schrimpf, M., Geiger, F., Cox, D., and DiCarlo, J. J. Simulating a primary visual cortex at the front of cnns improves robustness to image perturbations. In *Advances in Neural Information Processing Systems*, pp. 13073–13087, 2020.
- Deng, S. and Gu, S. Optimal conversion of conventional artificial neural networks to spiking neural networks. In *International Conference on Learning Representations*, pp. 1–14, 2021.
- Deng, S., Li, Y., Zhang, S., and Gu, S. Temporal efficient training of spiking neural network via gradient reweighting. In *International Conference on Learning Representations*, pp. 1–15, 2022.
- Diehl, P. U., Neil, D., Binas, J., Cook, M., Liu, S.-C., and Pfeiffer, M. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *International Joint Conference on Neural Networks*, pp. 1–8, 2015.
- Ding, J., Bu, T., Yu, Z., Huang, T., and Liu, J. K. SNN-RAT: Robustness-enhanced spiking neural network through regularized adversarial training. In *Advances in Neural Information Processing Systems*, pp. 1–14, 2022.
- Duan, C., Ding, J., Chen, S., Yu, Z., and Huang, T. Temporal effective batch normalization in spiking neural networks. In *Advances in Neural Information Processing Systems*, pp. 34377–34390, 2022.
- Esser, S. K., Merolla, P. A., Arthur, J. V., Cassidy, A. S., Appuswamy, R., Andreopoulos, A., Berg, D. J., McKinstry, J. L., Melano, T., Barch, D. R., di Nolfo, C., Datta, P., Amir, A., Taba, B., Flickner, M. D., and Modha, D. S. Convolutional networks for fast, energy-efficient neuromorphic computing. *The Proceedings of the National Academy of Sciences*, 113(41):11441–11446, 2016.
- Fang, W., Yu, Z., Chen, Y., Huang, T., Masquelier, T., and Tian, Y. Deep residual learning in spiking neural networks. In *Advances in Neural Information Processing Systems*, pp. 21056–21069, 2021a.
- Fang, W., Yu, Z., Chen, Y., Masquelier, T., Huang, T., and Tian, Y. Incorporating learnable membrane time constant to enhance learning of spiking neural networks. In *International Conference on Computer Vision*, pp. 2641–2651, 2021b.
- Finlay, C. and Oberman, A. M. Scaleable input gradient regularization for adversarial robustness. *Machine Learning with Applications*, 3:100017, 2021.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, pp. 1–11, 2015.
- Guo, Y., Chen, Y., Zhang, L., Liu, X., Wang, Y., Huang, X., and Ma, Z. IM-loss: Information maximization loss for spiking neural networks. In *Advances in Neural Information Processing Systems*, pp. 156–166, 2022.
- Han, B., Srinivasan, G., and Roy, K. RMP-SNN: Residual membrane potential neuron for enabling deeper high-accuracy and low-latency spiking neural network. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 13555–13564, 2020.

- Hao, Z., Bu, T., Shi, X., Huang, Z., Yu, Z., and Huang, T. Threaten spiking neural networks through combining rate and temporal information. In *International Conference on Learning Representations*, pp. 1–17, 2023a.
- Hao, Z., Ding, J., Bu, T., Huang, T., and Yu, Z. Bridging the gap between anns and snns by calibrating offset spikes. In *International Conference on Learning Representations*, 2023b.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Ho, J., Lee, B., and Kang, D. Attack-less adversarial training for a robust adversarial defense. *Applied Intelligence*, 52(4):4364–4381, 2022.
- Ho, N.-D. and Chang, I.-J. TCL: an ANN-to-SNN conversion with trainable clipping layers. In *Design Automation Conference*, pp. 793–798, 2021.
- Hu, Y., Tang, H., and Pan, G. Spiking deep residual networks. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8):5200–5205, 2023.
- Khan, M. A., El-Sayed, H., Malik, S., Zia, M. T., Khan, J., Alkaabi, N., and Ignatious, H. A. Level-5 autonomous driving - Are we there yet? A review of research literature. *ACM Computing Surveys*, 55(2):1–38, 2023.
- Kim, Y. and Panda, P. Revisiting batch normalization for training low-latency deep spiking neural networks from scratch. *Frontiers in Neuroscience*, 15:773954, 2021.
- Kim, Y., Park, H., Moitra, A., Bhattacharjee, A., Venkatesha, Y., and Panda, P. Rate coding or direct coding: Which one is better for accurate, robust, and energy-efficient spiking neural networks? In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 71–75, 2022.
- Kundu, S., Pedram, M., and Beerel, P. A. Hire-SNN: Harnessing the inherent robustness of energy-efficient deep spiking neural networks by training with crafted input noise. In *International Conference on Computer Vision*, pp. 5189–5198, 2021.
- Lee, C., Sarwar, S. S., Panda, P., Srinivasan, G., and Roy, K. Enabling spike-based backpropagation for training deep neural network architectures. *Frontiers in Neuroscience*, 14(119):1–22, 2020.
- Leontev, M., Antonov, D., and Sukhov, S. Robustness of spiking neural networks against adversarial attacks. In *International Conference on Information Technology and Nanotechnology*, pp. 1–6, 2021.
- Li, Y., Deng, S., Dong, X., Gong, R., and Gu, S. A free lunch from ANN: Towards efficient, accurate spiking neural networks calibration. In *International Conference on Machine Learning*, pp. 6316–6325, 2021.
- Liang, L., Hu, X., Deng, L., Wu, Y., Li, G., Ding, Y., Li, P., and Xie, Y. Exploring adversarial attack in spiking neural networks with spike-compatible gradient. *IEEE Transactions on Neural Networks and Learning Systems*, 34(5):2569–2583, 2021.
- Liang, L., Xu, K., Hu, X., Deng, L., and Xie, Y. Toward robust spiking neural network against adversarial perturbation. In *Advances in Neural Information Processing Systems*, pp. 10244–10256, 2022.
- Lipton, Z. C. The mythos of model interpretability. *Communications of the ACM*, 61(10):36–43, 2018.
- Liu, Y., Jiang, M., and Jiang, T. Transferable adversarial examples based on global smooth perturbations. *Computers & Security*, 121:102816, 2022.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. In *International Conference on Computer Vision*, pp. 9992–10002, 2021.
- Loshchilov, I. and Hutter, F. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, pp. 1–16, 2017.
- Maass, W. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10(9):1659–1671, 1997.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, pp. 1–23, 2018.
- Marchisio, A., Pira, G., Martina, M., Masera, G., and Shafique, M. DVS-Attacks: Adversarial attacks on dynamic vision sensors for spiking neural networks. In *International Joint Conference on Neural Networks*, pp. 1–9, 2021a.
- Marchisio, A., Pira, G., Martina, M., Masera, G., and Shafique, M. R-SNN: An analysis and design methodology for robustifying spiking neural networks against adversarial attacks through noise filters for dynamic vision sensors. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 6315–6321, 2021b.
- Mishra, S., Dash, A., and Jena, L. Use of deep learning for disease detection and diagnosis. In *Bio-inspired Neurocomputing*, pp. 181–201. Springer, Singapore, 2021.

- Mostafa, H. Supervised learning based on temporal coding in spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 29(7):3227–3235, 2017.
- Natarajan, B. K. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995.
- Neftci, E. O., Mostafa, H., and Zenke, F. 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, 2019.
- Nomura, O., Sakemi, Y., Hosomi, T., and Morie, T. Robustness of spiking neural networks based on time-to-first-spike encoding against adversarial attacks. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 69(9):3640–3644, 2022.
- Özdenizci, O. and Legenstein, R. Adversarially robust spiking neural networks through conversion. *arXiv preprint arXiv:2311.09266*, 2023.
- Ramirez, C., Kreinovich, V., and Argaez, M. Why ℓ_1 is a good approximation to ℓ_0 : A geometric explanation. *Journal of Uncertain Systems*, 7:203–207, 2013.
- Rathi, N. and Roy, K. DIET-SNN: A low-latency spiking neural network with direct input encoding and leakage and threshold optimization. *IEEE Transactions on Neural Networks and Learning Systems*, 34(6):3174–3182, 2021.
- Roy, K., Jaiswal, A., and Panda, P. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784):607–617, 2019.
- Sengupta, A., Ye, Y., Wang, R., Liu, C., and Roy, K. Going deeper in spiking neural networks: VGG and residual architectures. *Frontiers in Neuroscience*, 13(95):1–10, 2019.
- Shao, Z., Fang, X., Li, Y., Feng, C., Shen, J., and Xu, Q. EICIL: joint excitatory inhibitory cycle iteration learning for deep spiking neural networks. In *Advances in Neural Information Processing Systems*, pp. 1–12, 2023.
- Sharmin, S., Panda, P., Sarwar, S. S., Lee, C., Ponghiran, W., and Roy, K. A comprehensive analysis on adversarial robustness of spiking neural networks. In *International Joint Conference on Neural Networks*, pp. 1–8, 2019.
- Sharmin, S., Rathi, N., Panda, P., and Roy, K. Inherent adversarial robustness of deep spiking neural networks: Effects of discrete input encoding and non-linear activations. In *European Conference on Computer Vision*, pp. 399–414, 2020.
- Shen, J., Xu, Q., Liu, J. K., Wang, Y., Pan, G., and Tang, H. ESL-SNNs: An evolutionary structure learning strategy for spiking neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 86–93, 2023.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Stutz, D., Hein, M., and Schiele, B. Disentangling adversarial robustness and generalization. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6976–6987, 2019.
- Tanay, T. and Griffin, L. D. A boundary tilting perspective on the phenomenon of adversarial examples. *arXiv*, arXiv preprint arXiv:1608.07690, 2016.
- Tramèr, F., Kurakin, A., Papernot, N., Goodfellow, I., Boneh, D., and McDaniel, P. Ensemble adversarial training: Attacks and defenses. In *International Conference on Learning Representations*, pp. 1–20, 2018.
- Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., and Madry, A. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, pp. 1–23, 2019.
- Wong, E. and Kolter, J. Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, pp. 5283–5292, 2018.
- Wong, E., Rice, L., and Kolter, J. Z. Fast is better than free: Revisiting adversarial training. In *International Conference on Learning Representations*, pp. 1–12, 2020.
- Wu, Y., Deng, L., Li, G., Zhu, J., and Shi, L. Spatio-temporal backpropagation for training high-performance spiking neural networks. *Frontiers in Neuroscience*, 12(331):1–12, 2018.
- Wu, Y., Deng, L., Li, G., Zhu, J., Xie, Y., and Shi, L. Direct training for spiking neural networks: Faster, larger, better. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1311–1318, 2019.
- Xu, K., Shi, Z., Zhang, H., Wang, Y., Chang, K., Huang, M., Kailkhura, B., Lin, X., and Hsieh, C. Automatic perturbation analysis for scalable certified robustness and beyond. In *Advances in Neural Information Processing Systems*, pp. 1129–1141, 2020.
- Xu, N., Mahmood, K., Fang, H., Rathbun, E., Ding, C., and Wen, W. Securing the spike: On the transferability and security of spiking neural networks to adversarial examples. *arXiv preprint arXiv:2209.03358*, 2022a.

- Xu, Q., Li, Y., Shen, J., Zhang, P., Liu, J. K., Tang, H., and Pan, G. Hierarchical spiking-based model for efficient image classification with enhanced feature extraction and encoding. *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–9, 2022b.
- Yao, X., Li, F., Mo, Z., and Cheng, J. GLIF: A unified gated leaky integrate-and-fire neuron for spiking neural networks. In *Advances in Neural Information Processing Systems*, pp. 32160–32171, 2022.
- Yu, Z., Bu, T., Zhang, Y., Jia, S., Huang, T., and Liu, J. K. Robust decoding of rich dynamical visual scenes with retinal spikes. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. In *Proceedings of the British Machine Vision Conference*, pp. 1–15, 2016.
- Zenke, F., Bohté, S. M., Clopath, C., Comşa, I. M., Göltz, J., Maass, W., Masquelier, T., Naud, R., Neftci, E. O., Petrovici, M. A., et al. Visualizing a joint future of neuroscience and neuromorphic engineering. *Neuron*, 109(4): 571–575, 2021.
- Zhang, B., Cai, T., Lu, Z., He, D., and Wang, L. Towards certifying ℓ_∞ robustness using neural networks with ℓ_∞ -dist neurons. In *International Conference on Machine Learning*, pp. 12368–12379, 2021.
- Zhang, H., Chen, H., Xiao, C., Goyal, S., Stanforth, R., Li, B., Boning, D. S., and Hsieh, C. Towards stable and efficient training of verifiably robust neural networks. In *International Conference on Learning Representations*, pp. 1–25, 2020.
- Zhang, M., Wang, J., Wu, J., Belatreche, A., Amornpaisanon, B., Zhang, Z., Miriyala, V. P. K., Qu, H., Chua, Y., Carlson, T. E., et al. Rectified linear postsynaptic potential function for backpropagation in deep spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(5):1947–1958, 2022.
- Zhang, W. and Li, P. Temporal spike sequence learning via backpropagation for deep spiking neural networks. In *Advances in Neural Information Processing Systems*, pp. 12022–12033, 2020.
- Zheng, H., Wu, Y., Deng, L., Hu, Y., and Li, G. Going deeper with directly-trained larger spiking neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 11062–11070, 2021.
- Zhu, Y., Yu, Z., Fang, W., Xie, X., Huang, T., and Masquelier, T. Training spiking neural networks with event-driven backpropagation. In *Advances in Neural Information Processing Systems*, pp. 30528–30541, 2022.
- Zhu, Y., Fang, W., Xie, X., Huang, T., and Yu, Z. Exploring loss functions for time-based training strategy in spiking neural networks. *Advances in Neural Information Processing Systems*, pp. 65366–65379, 2023.

A. Proof of Theorem 4.3

Theorem 1. Suppose f represents an SNN and ϵ is the strength of an attack. Given an input image \mathbf{x} with corresponding label y , the ratio of $\rho_{adv}(f, \mathbf{x}, \epsilon, \ell_\infty)$ and $\rho_{rand}(f, \mathbf{x}, \epsilon, \ell_\infty)$ is upper bounded by the sparsity of $\nabla_{\mathbf{x}} f_y$:

$$3 \leq \frac{\rho_{adv}(f, \mathbf{x}, \epsilon, \ell_\infty)}{\rho_{rand}(f, \mathbf{x}, \epsilon, \ell_\infty)} \leq 3 \|\nabla_{\mathbf{x}} f_y(\mathbf{x})\|_0. \quad (16)$$

Proof. We assume f to be differentiable, where the surrogate gradient is used. When ϵ is small, we can expand $f_y(\mathbf{x} + \epsilon \cdot \delta)$ at $f(\mathbf{x})$ by the first-order Taylor expansion

$$f_y(\mathbf{x} + \epsilon \cdot \delta) \approx f_y(\mathbf{x}) + \epsilon \nabla f_y(\mathbf{x})^T \delta. \quad (17)$$

So, $f_y(\mathbf{x} + \epsilon \cdot \delta) - f_y(\mathbf{x}) \approx \epsilon \nabla f(\mathbf{x})^T \delta$. As $\delta \in \mathbb{R}^m$ and $\delta_i \sim \text{Unif}([-1, 1])$, we have

$$\mathbb{E}(\delta_i \delta_j) = \begin{cases} 0 & i \neq j \\ \frac{1}{3} & i = j. \end{cases} \quad (18)$$

Therefore, $\rho_{rand}(f, \mathbf{x}, \epsilon, \ell_\infty)$ can be approximated as follows:

$$\rho_{rand}(f, \mathbf{x}, \epsilon, \ell_\infty) = \mathbb{E}_{\delta \sim \text{Unif}(cube)} (f_y(\mathbf{x} + \epsilon \cdot \delta) - f_y(\mathbf{x}))^2 \approx \epsilon^2 \nabla f_y(\mathbf{x})^T \mathbb{E}_\delta (\delta \delta^T) \nabla f_y(\mathbf{x}) = \frac{1}{3} \epsilon^2 \|\nabla f(\mathbf{x})\|_2^2. \quad (19)$$

On the other hand,

$$\begin{aligned} \rho_{adv}(f, \mathbf{x}, \epsilon, \ell_\infty) &= \sup_{\delta \sim \text{Unif}(cube)} (f_y(\mathbf{x} + \epsilon \cdot \delta) - f_y(\mathbf{x}))^2 \approx \epsilon^2 \left(\sup_{\delta \sim \text{Unif}(cube)} |\nabla f_y(\mathbf{x})^T \delta| \right)^2 \\ &= \epsilon^2 (\nabla f_y(\mathbf{x})^T \text{sign}(\nabla f_y(\mathbf{x})))^2 = \epsilon^2 \|\nabla f_y(\mathbf{x})\|_1^2. \end{aligned} \quad (20)$$

Consequently, the gap between $\rho_{adv}(f, \mathbf{x}, \epsilon, \ell_\infty)$ and $\rho_{rand}(f, \mathbf{x}, \epsilon, \ell_\infty)$ can be measured by

$$\frac{\rho_{adv}(f, \mathbf{x}, \epsilon, \ell_\infty)}{\rho_{rand}(f, \mathbf{x}, \epsilon, \ell_\infty)} \approx 3 \frac{\|\nabla f_y(\mathbf{x})\|_1^2}{\|\nabla f_y(\mathbf{x})\|_2^2}, \quad (21)$$

which can be bounded by

$$3 \leq \frac{\rho_{adv}(f, \mathbf{x}, \epsilon, \ell_\infty)}{\rho_{rand}(f, \mathbf{x}, \epsilon, \ell_\infty)} \leq 3 \|\nabla f_y(\mathbf{x})\|_0. \quad (22)$$

[proof of the inequality] For an m -dimensional vector $u \in \mathbb{R}^m$, we have $\|u\|_1 \geq \|u\|_2$. Because

$$\|u\|_1^2 = \left(\sum_{i=1}^m |u_i| \right)^2 = \sum_{i=1}^m u_i^2 + \sum_i \sum_{j \neq i} |u_i u_j| \geq \sum_{i=1}^m u_i^2 = \|u\|_2^2. \quad (23)$$

Moreover, let $a \in \mathbb{R}^m$ be an m -dimensional vector with $a_i = \text{sign}(u_i)$. Then

$$\|u\|_1 = \sum_{i=1}^m |u_i| = \sum_{i=1}^m u_i a_i \leq \left(\sum_{i=1}^m u_i^2 \right)^{\frac{1}{2}} \left(\sum_{i=1}^m a_i^2 \right)^{\frac{1}{2}} \quad (\text{Cauchy Schwartz inequality}) = \|u\|_2 \sqrt{\|u\|_0} \quad (24)$$

□

B. Derivation of Equation (12)

Let \mathbf{x} denote the image, and $\{\mathbf{x}[1], \mathbf{x}[2], \dots, \mathbf{x}[T]\}$ represent the input image series. In our paper, we use $\mathbf{x}[t] = \mathbf{x}$ for all $t = 1, \dots, T$. The network is denoted by f and the output of the network f in the last layer is a vector $f^L(\mathbf{x}) \in \mathbb{R}^{N \times 1}$, where N represents the number of classes, i.e.

$$f^L(\mathbf{x}) = \left(\sum_{t=1}^T s_1^L[t], \dots, \sum_{t=1}^T s_N^L[t] \right). \quad (25)$$

Based on Equation (9), the component related to the true label y is defined as

$$f_y(\mathbf{x}) = \frac{e^{f_y^L}}{e^{f_y^L} + e^{f_{\tilde{y}}^L}} = \frac{e^{\sum_{t=1}^T s_y^L(t)}}{e^{\sum_{t=1}^T s_y^L(t)} + e^{\sum_{t=1}^T s_{\tilde{y}}^L(t)}}, \quad (26)$$

According to the chain rule, the gradient of $f_y(\mathbf{x})$ with respect to the input \mathbf{x} is

$$\nabla_{\mathbf{x}} f_y(\mathbf{x}) = \frac{\partial f_y}{\partial f_y^L} \cdot \nabla_{\mathbf{x}} \left(\sum_t s_y^L[t] \right) + \frac{\partial f_y}{\partial f_{\tilde{y}}^L} \cdot \nabla_{\mathbf{x}} \left(\sum_t s_{\tilde{y}}^L[t] \right). \quad (27)$$

In this formula, the gradient of $\nabla_{\mathbf{x}} \left(\sum_t s_i^L[t] \right)$ ($i = y, \tilde{y}$) in the last layer can be further express as

$$\nabla_{\mathbf{x}} \left(\sum_t s_i^L[t] \right) = \sum_t \nabla_{\mathbf{x}} s_i^L[t] = \sum_t \sum_{\tilde{t}=1}^t \nabla_{\mathbf{x}[\tilde{t}]} s_i^L[t]. \quad (28)$$

Finally, the gradient $\nabla_{\mathbf{x}} f_y(\mathbf{x})$ is written as

$$\nabla_{\mathbf{x}} f_y(\mathbf{x}) = \sum_{i=y, \tilde{y}} \left(\frac{\partial f_y}{\partial f_i^L} \left(\sum_{t=1}^T \sum_{\tilde{t}=1}^t \nabla_{\mathbf{x}[\tilde{t}]} s_i^L[t] \right) \right). \quad (29)$$

C. Proof of Proposition 4.4

Proposition 1. Let d denote the signed input gradient direction: $d = \text{sign}(\nabla_{\mathbf{x}} f_y(\mathbf{x}))$, and h be the finite difference step size. Then, the ℓ_1 gradient norm can be approximated as:

$$\|\nabla_{\mathbf{x}} f_y(\mathbf{x})\|_1 \approx \left| \frac{f_y(\mathbf{x} + h \cdot d) - f_y(\mathbf{x})}{h} \right| \quad (30)$$

Proof. The first order Taylor estimation of $f_y(\mathbf{x} + h \cdot d)$ at point \mathbf{x} is

$$f_y(\mathbf{x} + h \cdot d) \approx f_y(\mathbf{x}) + h \cdot \nabla_{\mathbf{x}} f_y^T(\mathbf{x}) d = f_y(\mathbf{x}) + h \|\nabla_{\mathbf{x}} f_y(\mathbf{x})\|_1. \quad (31)$$

Therefore, $\|\nabla_{\mathbf{x}} f_y(\mathbf{x})\|_1$ can be approximated by

$$\left| \frac{f_y(\mathbf{x} + h \cdot d) - f_y(\mathbf{x})}{h} \right| \quad (32)$$

□

D. Training Settings

We use the same training settings for all architectures and datasets. Our data augmentation techniques include RandomCrop, RandomHorizontalFlip, and zero-mean normalization. During training, we use the CrossEntropy loss function and Stochastic Gradient Descent optimizer with momentum. The learning rate η is controlled by the cosine annealing strategy (Loshchilov & Hutter, 2017). We utilize the Backpropagation Through Time (BPTT) algorithm with a triangle-shaped surrogate function, as introduced by (Esser et al., 2016). When incorporating sparsity gradient regularization, we set the step size of the finite difference method to 0.01. Also, we use a $\lambda = 0.002$ on CIFAR-10/CIFAR10-DVS and $\lambda = 0.001$ on CIFAR-100 for SR* method. For vanilla SR, we set $\lambda = 0.008$ on CIFAR-10 and $\lambda = 0.002$ on CIFAR-100/CIFAR10-DVS. The detailed training hyper-parameters are listed below.

Table 5: Detailed training setting.

Timestep	Initial LR	Batchsize	Weight Decay	Epochs	Momentum	h	AT	ϵ	PGD-step
CIFAR-10/100 Dataset									
8	0.1	64	5e-4	200	0.9	0.01	PGD5	2/255	0.01
CIFAR10-DVS Dataset									
10	0.1	24	5e-4	200	0.9	0.01	FGSM	2/255	/

E. Evaluation Settings

As mentioned in the main text, we consider an ensemble attack approach for SNNs. This involves utilizing a diverse set of surrogate gradients and considering both STBP-based and RGA-based attacks. We conduct the following attacks as the ensemble attack. We consider an ensemble attack to be successful for a test sample as long as the model is fooled with any of the attacks from the ensemble.

- **STBP-based attack with triangle-shaped surrogate function**, with the hyper-parameter $\gamma = 1$ (Esser et al., 2016).

$$\frac{\partial s_i^l(t)}{\partial u_i^l(t)} = \frac{1}{\gamma} \|\gamma - |u_i^l(t) - \theta|\|. \quad (33)$$

- **STBP-based attack with sigmoid-shaped surrogate function**, with the hyper-parameter γ being 4.

$$\frac{\partial s_i^l(t)}{\partial u_i^l(t)} = \frac{1}{1 + \exp(-\gamma(u_i^l(t) - \theta))} \quad (34)$$

- **STBP-based attack with arc tangent surrogate function**, with the hyper-parameter γ set to 2.

$$\frac{\partial s_i^l(t)}{\partial u_i^l(t)} = \frac{\gamma}{2(1 + (\frac{\pi}{2}\gamma(u_i^l(t) - \theta))^2)} \quad (35)$$

- **RGA-based attack with rate-based gradient estimation**, where the setting follows the paper (Bu et al., 2023).

F. Coefficient Parameter Search on CIFAR-100

Figure 6 (left) shows the relationship between the coefficient parameter λ , clean accuracy, and adversarial accuracy on the CIFAR-100 dataset. To make a trade-off between classification accuracy on clean images and adversarial images, we choose

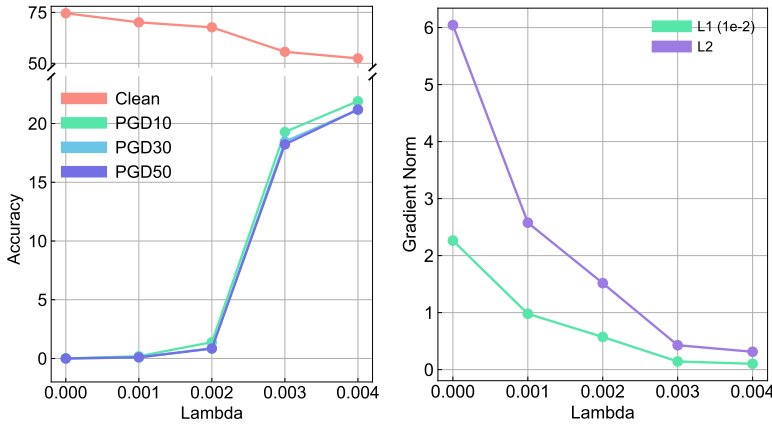


Figure 6: The influence of the coefficient parameter λ on classification accuracy and gradient sparsity. Left: Fluctuations in clean accuracy and adversarial accuracy under PGD attacks. Right: The ℓ_1 and ℓ_2 norms of the gradient with varying λ .

$\lambda = 0.002$ for SR-WRN-16 on the CIFAR-100 dataset. Figure 6 (right) illustrates that the ℓ_1 and ℓ_2 norm of $\nabla_x f_y$ decrease significantly with the increase of λ .

G. Extensive Evaluations With Varying Widths of Surrogate Functions

For the ensemble attack to be meaningful and reveal any impact of gradient obfuscation, we run an evaluation with different widths γ in the surrogate function. Specifically, we apply the attack with $\gamma \in [0.1, 3.0]$ in fine-grained steps of 0.1, and report the results of the PGD10 attack on VGG-11 models with different training algorithms on the CIFAR-10 dataset in Table 6.

Table 6: The classification accuracy (%) under the ensemble attack with different γ .

Attacks	RAT	AT	SR*
PGD10 (w/o ensemble)	16.16	21.32	33.67
PGD10 (w/ ensemble)	11.53	18.18	30.54
PGD10 ($\gamma \in [0.1, 3.0]$)	11.87	16.16	27.06

We compare three different attack combinations to evaluate the impact of different surrogate functions on the attack strength. We select RAT, PGD5-AT, and SR*(PGD5+SR) models as the target models. For PGD10(w/o ensemble), we only use the Triangle-shaped surrogate function, which is identical to the one used in training. For PGD10 (w/ ensemble), we use the attack combination as described in Section 5. For PGD10($\gamma \in [0.1, 3.0]$), we incorporate 30 different Triangle-shaped surrogate functions with widths ranging from $[0.1, 30]$.

We find that both ensemble attack methods significantly improve attack performance and mitigate the impact of gradient obfuscation. This indicates that both the shape and width of the surrogate function can influence the capability of the adversary. Although the PGD10($\gamma \in [0.1, 3.0]$) attack is slightly more effective than the ensemble attack used in Section 5, it uses a 30-fold fine-grained grid search over attack hyperparameters for each image, which is considerably more computationally expensive.

In conclusion, we demonstrate that changing the width of surrogate functions does not significantly influence the capability of the adversary any better than using different surrogate gradient shapes. Additionally, the proposed SR* strategy exhibits improved robustness under both scenarios.

H. Comparison in Computational Cost

The computational costs for one epoch training of various algorithms, including vanilla, PGD5 AT, RAT, SR, and SR*(PGD5+SR), on the CIFAR-10 dataset using the VGG11 architecture are summarized in Table 7. From the table, we observe that a single SR incurs a computational cost 1.5 times that of RAT but consumes less than half the time needed by PGD5 AT. The computational cost of SR* is the highest among all training algorithms since it combines both SR and AT. However, models trained with SR* achieve the best robustness compared to models trained with other methods.

Table 7: The computational cost in one epoch of different training algorithms.

Vanilla	PGD5-AT	RAT	SR	SR*
65s	459s	134s	193s	583s