

# Security-Aware Approximate Spiking Neural Networks

Syed Tihaam Ahmad, Ayesha Siddique, Khaza Anuarul Hoque

*Department of Electrical Engineering and Computer Science*

*University of Missouri, Columbia, MO, USA*

tawm9@umsystem.edu, ayesha.siddique@mail.missouri.edu, hoquek@missouri.edu

**Abstract**—Deep Neural Networks (DNNs) and Spiking Neural Networks (SNNs) are both known for their susceptibility to adversarial attacks. Therefore, researchers in the recent past have extensively studied the robustness and defense of DNNs and SNNs under adversarial attacks. Compared to accurate SNNs (AccSNN), approximate SNNs (AxSNNs) are known to be up to 4X more energy-efficient for ultra-low power applications. Unfortunately, the robustness of AxSNNs under adversarial attacks is yet unexplored. In this paper, we first extensively analyze the robustness of AxSNNs with different structural parameters and approximation levels under two gradient-based and two neuromorphic attacks. Then, we propose two novel defense methods, i.e., precision scaling and approximate quantization-aware filtering (AQF), for securing AxSNNs. We evaluated the effectiveness of these two defense methods using both static and neuromorphic datasets. Our results demonstrate that AxSNNs are more prone to adversarial attacks than AccSNNs, but precision scaling and AQF significantly improve the robustness of AxSNNs. For instance, a PGD attack on AxSNN results in a 72% accuracy loss compared to AccSNN without any attack, whereas the same attack on the precision-scaled AxSNN leads to only a 17% accuracy loss in the static MNIST dataset (4X robustness improvement). Similarly, a Sparse Attack on AxSNN leads to a 77% accuracy loss when compared to AccSNN without any attack, whereas the same attack on an AxSNN with AQF leads to only a 2% accuracy loss in the neuromorphic DVS128 Gesture dataset (38X robustness improvement).

**Index Terms**—Spiking Neural Networks, Approximate Spiking Neural Networks, Adversarial Robustness, Approximate Defense.

## I. INTRODUCTION

Spiking neural networks (SNNs) are the third generation of neural networks that employ event-driven computing capabilities [1]. In recent years, many SNN models of different sizes have been developed for data analytics, such as gesture recognition, object detection, and image classification. However, large-sized SNN models are known for recognizing more features than small ones. Consequently, the state-of-the-art large-sized SNN models have numerous parameters that need to be considered in both the training and inference phases. This limits the deployment of SNNs on ultra-low power resource-constrained edge devices. To handle this problem, approximate computing in SNNs has recently emerged as an energy-efficient solution. Approximate computing-based SNNs (AxSNNs) relax the abstraction of near-perfect accuracy in error-resilient applications for their low energy consumption. For instance, AxSNNs obtained via

approximating the weights can reduce the energy consumption by 4X [2] compared to the accurate SNNs (AccSNNs).

Similar to traditional deep neural networks (DNNs), AxSNNs are also prone to adversarial attacks. Adversarial attacks are known for being very stealthy as they add minimal perturbation noise to the inputs, which are imperceptible to the human eye, yet successfully fool the SNN classifiers. [3]. Motivated by this, the robustness analysis and adversarial defense in AccSNNs have been thoroughly investigated in several recent works [4]–[7]. Very recently, in [8], the authors showed that approximate DNNs (AxDNNs) are more prone to adversarial attacks as compared to accurate DNNs (AccDNN). However, surprisingly, the robustness of AxSNNs is yet unexplored. Therefore, a more comprehensive study is required to understand the inherent behavior of AxSNNs vs. AccSNNs, especially under adversarial attacks. Exploration of such behaviors can enable designing defense techniques tailored specifically for AxSNNs.

### A. Motivational Case Study and Key Observations

As a preliminary study for motivating our research, we conducted a motivational case study highlighting the impact of adversarial attacks on AccSNNs vs. AxSNNs. For this purpose, we first trained a 5-layered AccSNN, having 3 convolutional layers and 2 fully-connected layers for classifying the MNIST [9] dataset. Then, we built an AxSNN (using approximation level 0.1) as an approximate counterpart of the AccSNN. Finally, we compared the performance of AccSNN and AxSNN under the  $l_\infty$  norm-based projected gradient descent (PGD) attack by varying the perturbation budget  $\epsilon$  ranging from 0 to 1.0. The results are presented in Fig. 1. We observe that the AxSNN is significantly less robust when compared to AccSNNs under attack. For instance, when there is no attack ( $\epsilon=0$ ), the accuracy of AccSNN and AxSNN is 97% and 52%, respectively. However, for ( $\epsilon=0.5$ ), the accuracy of

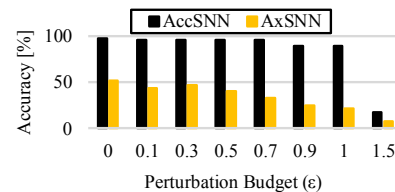


Figure 1: Robustness comparison of AccSNN and AxSNN under PGD attack with different perturbation budgets.

AccSNN and AxSNN is 95% and 40%, respectively, which shows a 55% difference in their accuracy. Furthermore, when the perturbation budget was varied to 1.0 ( $\epsilon=1.0$ ), we can observe a 68% difference while comparing the accuracy of the AccSNN and AxSNN. These outcomes motivated us to thoroughly investigate the robustness of the AxSNNs and explore potential defense techniques to design robust AxSNNs.

### B. Novel Contributions

In this paper, we present a security-aware AxSNNs design method with the following novel contributions:

- 1) A novel design approach for designing adversarially robust AxSNNs by identifying their robustness-aware knobs through *precision scaling*, i.e., through finding the appropriate combination approximation levels, structural parameters (threshold voltage, time steps), and quantization. We also propose a defense method through Approximate Quantization-aware filtering (AQF), specifically effective against neuromorphic attacks. [Section IV]
- 2) An extensive vulnerability analysis of AxSNNs vs. AccSNNs against two gradient-based attacks and two neuromorphic attacks under different threshold voltage, time steps, precision scales, and approximation levels. Specifically, we evaluated the impact of Projected Gradient Descent (PGD) and Basic Iterative Method (BIM) attacks on AccSNN and AxSNN with a static dataset MNIST [9]. On the contrary, we evaluated the impact of Sparse and Frame attacks on AccSNN and AxSNN with the neuromorphic dataset DVS128 Gesture [10]). [Section V]

Our results demonstrate that AxSNNs are more prone to adversarial attacks than AccSNNs. However, the precision scaling and approximate quantization aware filtering improve their robustness significantly. For instance, a PGD attack with perturbation budget 1.0 on AxSNNs results in a 72% accuracy loss, whereas the same attack on AccSNNs results in only a 9% accuracy loss in MNIST classification. Interestingly, the same AxSNN with precision scaling shows only a 17% accuracy loss, indicating a 4X robustness improvement. Similarly, a Sparse Attack on AxSNN leads to a 77% accuracy loss when compared to AccSNN without any attack in the DVS128 Gesture classification. However, after using our designed approximate quantization aware filter the accuracy loss is just 2%, indicating a 38X improvement in robustness.

## II. PRELIMINARIES

This section provides a brief overview of SNNs and adversarial attacks to understand the paper better.

**Approximate Spiking Neural Networks:** AxSNNs employ approximate computing to trade their classification accuracy with energy efficiency in ultra-low power applications. AxSNNs typically associate an approximation level  $a_{th}$  with each spiking neuron. The  $a_{th}$  determines if the respective neurons should be activated or deactivated based on the sensitivity of the neurons to errors and spiking activity [11]. Similar to AccSNNs, AxSNNs use the

standard leaky-integrate-and-fire (LIF) model where, when the membrane potential exceeds the threshold voltage, the neuron emits an output spike and resets its membrane potential. Specifically, they process the spike-encoded inputs which are most commonly encoded using rate encoding, where the activation activity corresponds to the mean firing rates of spikes over certain time steps. The time steps refer to the observation period when the SNN receives the same input.

**Adversarial Attacks:** The adversarial attacks are small perturbations that cause the classifier to predict the false labels. Examples of such attacks include gradient-based PGD and BIM, which are strong attacks in the adversarial machine learning domain. The gradient-based attacks exploit the concept of back-propagation; however, instead of calculating the gradient with respect to the weights of the model, they craft the adversarial examples by perturbing the input images. Recent studies show that these attacks cannot be used for perturbing the neuromorphic datasets due to their event-driven nature. Therefore, specialized neuromorphic attacks, such as Sparse and Frame Attacks are used [6]. A Sparse Attack is a stealthy attack that iteratively perturbs the neuromorphic images based on the loss function of output label probability to generate perturbed events. On the other hand, a Frame Attack is a simple yet effective neuromorphic attack that generates perturbed events by attacking every pixel of the boundary for all the events.

## III. THREAT MODEL

In this section, we present a threat model for exploring the adversarial robustness of AxSNNs.

**Adversary's Knowledge:** We assume that the adversary uses an accurate classifier model for crafting the adversarial examples. Furthermore, the adversary has partial knowledge about AxSNN, i.e., the internal architecture of the classifier model is known, but the inexactness and model parameters such as threshold voltage and time-steps, precision scale, and approximation levels are not known to the adversary.

**Attack Generation:** The adversary is assumed capable of evading the classifier model by tampering the input images in the prediction phase without influencing the training data. The adversary crafts the adversarial examples by finding the perturbations that maximize the loss of a model on input while keeping the perturbation magnitude lower than the perturbation budget  $\epsilon$ . As mentioned earlier, we use iterative gradient-based attacks specifically,  $l_\infty$  norm-based *BIM and PGD*, which are considered high-strength attacks for the static datasets. We also employ neuromorphic attacks specifically, sparse and frame attacks which are stealthy yet effective in perturbing the neuromorphic images with high resolutions. We used  $\epsilon=1.0$  in this paper because the accuracy of both AccSNN and AxSNN drops significantly after this value, so it becomes non-recoverable. As an example, it drops to 10% with  $\epsilon=1.5$ .

## IV. SECURITY-AWARE AXSNN DESIGN APPROACH

In this section, we discuss our proposed approach for designing a robust and secure AxSNN in detail.

### A. Precision-scaling

Traditionally, the approximation levels are determined by identifying the maximum tolerable perturbation (MTP) in a neural network [12]. However, this becomes challenging in precision-scaled AxSNNs due to the variation in accuracy with the change in the threshold voltage, time steps and precision scales. Intuitively, an increase in the number of time steps and threshold voltage may lead to a higher number of insignificant spikes by some neurons and hence, affect the classification accuracy. Skipping such neurons has the potential to improve the accuracy of AxSNNs; however, their robustness can decrease under attacks. Since precision scaling has the potential to improve the robustness of AccDNNs [12], determining approximation on the basis of the precision scales in AxSNNs can improve the adversarial robustness. However, exploring a precision scale in addition to threshold voltage and time steps for robust approximation seems challenging. In this paper, we determine the approximation levels  $a_{th}$  in AxSNNs by using the following equation:

$$a_{th} = (cN_s/T) \cdot \min(1, V_m/V_{th}) \cdot \sum_{i=1}^c w_i^p, \quad (1)$$

where  $c$ ,  $N_s$ ,  $T$ ,  $V_m$ ,  $V_{th}$  and  $w^p$  denote the number of connections to output, number of spikes, time steps, membrane potential, threshold voltage, and precision scaled weight of neuron. Furthermore,  $\min(1, V_m/V_{th})$  is the spike probability. The maximum spike probability is 1 when  $V_m$  crosses  $V_{th}$  otherwise,  $V_m/V_{th}$ . The spike probability is weighted by the mean of all weights corresponding to a connection  $c$  i.e.,  $\sum_{i=1}^c w_i^p$  which includes the precision scaling of weights.

Algorithm 1 delineates the steps involved in implementing this equation for robustness in AxSNNs. We first initialize a counter  $adv$  for successful attack generation (Line 1). Then, we train an AccSNN model with the given threshold voltage  $v$  and time steps  $t$  and save the trained model and the corresponding weights of each layer  $l$  (Line 3). This learning phase is quantitatively verified with a quality constraint  $Q$  i.e., minimum baseline accuracy below which we consider SNN learning inefficient (Line 4). The value of  $Q$  depends on the given SNN architecture, dataset, and application. Then, we craft the adversarial examples for an adversarial attack and perturbation budget  $\epsilon$  (Line 5). The adversarial defense through precision scaling starts from Line 6 where the trained weights are sorted initially in ascending order. This sorting helps us in approximating the weights according to their significance later on. Afterwards, we perform precision scaling on the basis of each precision scale  $s$  and calculate the mean of all connections in layer  $l$  (Lines 8-9). Using this mean, we determine  $a_{th}$  (as discussed earlier) for each layer  $l$  and approximate the precision-scaled model by removing the connections having weights below  $a_{th}$  (Line 10). If a neuromorphic dataset is given as an input and the corresponding flag  $F_d$  is high then, we also use a special approximate quantization aware mask filter (AQF) from Algorithm 2 as discussed in Section IV-B (Line 12-14). Next, the algorithm checks if the crafted adversarial example

can fool AxSNN, i.e., if the attack succeeds in forcing the output to a wrong label, and accordingly increment the adversarial success counter (Line 15-18). Lastly, we evaluate the robustness for the perturbation budget  $\epsilon$  as the rate of attacks for which the adversary failed to generate an adversarial example that fools the victim SNN (Line 21). In this paper, we use this algorithm to find the approximation level  $a_{th}$  with the robust set of threshold voltage, time steps, and precision scales. Therefore, we compare the accuracies of the precision-scaled AxSNN models across all these parameters and return their values which meet our quality constraint (Lines 22-24).

---

#### Algorithm 1: Precision-Scaling in AxSNNs

---

**Inputs :** Type of adversarial attack:  $attack$ ;  
 Perturbation budget:  $\epsilon$ ;  
 Time steps:  $T = [t_1, t_2, \dots, t_n]$ ;  
 Threshold Voltage:  $V_{th} = [v_1, v_2, \dots, v_n]$ ;  
 Train dataset:  $\mathcal{D}_{tr} = (X, L)$ ;  
 Test dataset:  $\mathcal{D}_{ts} = (x, l)$ ;  
 Perturbation budget:  $\epsilon$ ;  
 Precision-scaling level:  $s_l = [s_1, s_2, \dots, s_n]$ ;  
 Quality constraint:  $Q$ ;  
 Neuromorphic Dataset Flag:  $F_d$

**Outputs:** Robustness level  $R$ , Best  $V_{th}$ , time steps  $ts$ , approximation level  $a_{th}$  and precision-scaling level  $s$

```

1:  $adv = 0$ ;
2: for each  $(v, ts)$  in  $(V_{th}, T)$  do
3:    $(model, w_l) = \text{trainAccurateSNN}(v, ts, \mathcal{D}_{tr})$ 
4:   if  $\text{Accuracy}(model) \geq Q$  then
5:      $(x_k^*, l_k^*) = \text{AdvExGen}(model, \epsilon, attack, X_k^t)$ 
6:      $w = \text{SortInAscendingOrder}(w_l)$ ;
7:     for each  $s$  in  $s_l$  do
8:        $w^p = \text{PrecisionScaling}(w, s)$ ;
9:        $m_l^c = \sum_{j=1}^c (w_j^p)$ ;
10:       $a_{th} = (cN_s/T) m_l^c \cdot \min(1, V_m / v)$ 
11:       $model = \text{ApproximateSNN}(model, w^p, m_l^c, a_{th})$ ;
12:      if  $T_d$  is TRUE then
13:         $\mathcal{D}_{ts} = \text{ApproximateQuantizedFilter}(qt, \mathcal{D}_{ts})$ 
14:      end if
15:       $(x_k', l_k') = \text{AdvAttacks}(model, \epsilon, attack, x_k^*, l_k^*)$ 
16:      if  $l_k' \neq l_k$  then
17:         $adv++$ ;
18:      else
19:        NOP;
20:      end if
21:       $R(\epsilon_i) = (1 - adv / \text{size}(\mathcal{D}_{ts})) * 100$ ;
22:      if  $R \geq Q$  then
23:        return  $(R, V_{th}, T, s, a_{th})$ 
24:      end if
25:    end for
26:  end if
27: end for
```

---

### B. Approximate Quantization Aware Filtering

For SNNs feeded by dynamic vision sensors (DVS), the above-discussed defense techniques for frame-based sensors cannot be directly applied due to the event-driven nature of neuromorphic images. Therefore, we present an additional approximate quantization aware filter (AQF) to remove uncorrelated events from the neuromorphic images. Our proposed Algorithm 2 removes adversarial perturbation noise from event-driven neuromorphic dataset  $E$ . The neuromorphic dataset  $E$  is represented in the form of  $(x, y, p, t)$ , where  $x$ ,  $y$ ,  $p$  and  $t$  denote the x-coordinate, the y-coordinate, the polarity

and the timestamp of the event  $E$  respectively. The events  $e$  are associated with a spatio-temporal domain and hence, they are correlated. Our algorithm calculates the correlation between events  $e$ . If the correlation for events  $e$  is lower than certain spatial-temporal thresholds ( $s, T1, T2$ ) then, they are removed because these events with very low correlation are more likely to be noisy due to adversarial perturbations.

---

**Algorithm 2:** ApproximateQuantizedFilter

---

**Inputs :** List of events:  $D_{ts}(x, y, p, t) = Events$   
Quantization step:  $q_t$ ;  
**Outputs:** Filtered Quantized dataset  $D_q$

```

1:  $M = 0$ 
2:  $activity = 0, s = 2, T1 = 5, T2 = 50$ 
3: for  $e$  in  $Events$  do
4:    $e = \text{round}(e/q_t) \cdot q_t$ 
5:   for  $i$  in  $(x_e - s, x_e + s)$  do
6:     for  $j$  in  $(y_e - s, y_e + s)$  do
7:       if  $\text{not}(i == x_e \text{ and } j == y_e)$  then
8:          $M[i][j] = t_e$ 
9:       end if
10:      if  $\text{not}(i == x_e \text{ and } j == y_e)$  then
11:         $activity[i][j] += 1$ 
12:      end if
13:    end for
14:  end for
15:  if  $activity[i][j] > T1$  then
16:     $M[i][j] = 1$ 
17:  end if
18:  if  $t_e - M[x_e][y_e] > T2$  or  $M[x_e][y_e] == 1$  then
19:    Remove  $e$  from  $Events$ 
20:  end if
21: end for
22:  $D_q = Events$ 
23: return  $D_q$ 

```

---

Algorithm 2 delineates the steps involved in approximate quantization aware filtering of perturbed events in the neuromorphic dataset. First, the dataset is quantized with a fixed quantization step  $q_t$  (Line 4) then, the uncorrelated values are checked in each event  $e$  of the dataset (Lines 5-9). With each low correlation, a counter variable *activity* is increased and flagged (Lines 10-16). Finally, the flagged uncorrelated values events  $e$  are removed from the dataset (Lines 18-20) to get a quantized and filtered dataset  $D_q$  (Line 23).

## V. RESULTS AND DISCUSSIONS

In this section, we present our results for the adversarial vulnerability analysis and adversarial defense in AxSNNs.

### A. Datasets and Architectures

We use both static and neuromorphic datasets i.e., MNIST [9] and DVS128 Gesture [10]. Both of these datasets are common for evaluating the performance of SNNs [13] in embedded platforms. As a classifier for the MNIST dataset, we employed a 7-layered SNN with three convolutional layers, two pooling, and two fully connected layers. The test accuracy of this architecture on clean inputs (without any adversarial attack) is 97%. On the other hand, as a classifier for the DVS128 Gesture dataset, we employed an 8-layered SNN with two convolutional and two fully connected layers, three pooling, and one dropout layer. The test accuracy of this architecture on clean inputs is 92%.

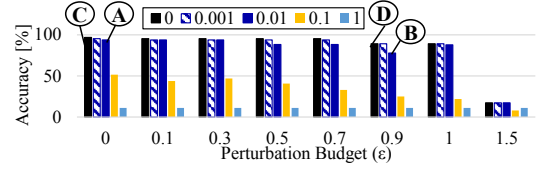


Figure 2: Robustness analysis of AccSNN (approximation level 0) and AxSNN MNIST classifier under PGD attack for approximation levels 0.001, 0.01, 0.1 and 1.

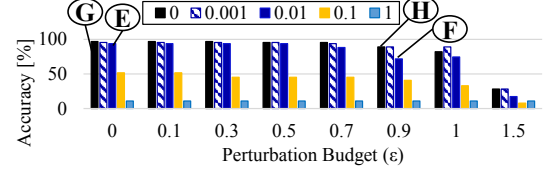


Figure 3: Robustness analysis of AccSNN (approximation level 0) and AxSNN MNIST classifier under BIM attack for approximation levels 0.001, 0.01, 0.1 and 1.

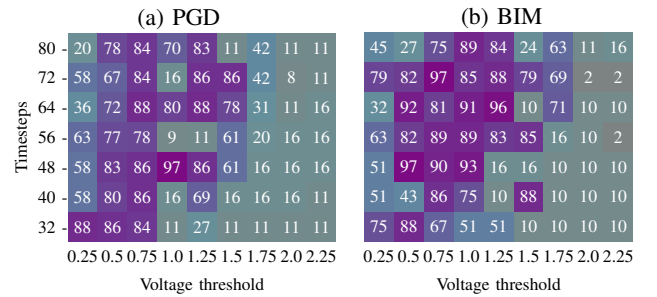


Figure 4: Accuracy of AxSNN (approx. level 0.01, precision-scale FP32) under attack ( $\epsilon = 1$ ) for MNIST.

1) Vulnerability analysis of AxSNNs: We first evaluate the adversarial robustness of AxSNNs by varying the approximation levels and comparing them with AccSNNs under different perturbation budgets for the static MNIST dataset. We keep threshold voltage and time steps constant at 0.25 and 32 in this experiment. Our results in Fig. 2 and Fig. 3 show that the robustness of AxSNNs decreases with an increase in the perturbation budget. For example, a PGD attack ( $\epsilon = 0.9$ ) drops the accuracy of AxSNN with approximation level 0.01 from 93% ( $\epsilon = 0$ ) to 77% (see labels A and B in Fig. 2) whereas the same attack drops the accuracy of the AccSNN from 96% to 89% see Label C and D in Fig. 2. This indicates a 7% accuracy loss for the AccSNN under PGD attack, whereas the same attack causes accuracy loss of 16% for the AxSNN. In this case, the AxSNN is 2X more vulnerable than AccSNN. Likewise, a BIM attack ( $\epsilon = 0.9$ ) drops the accuracy of the AxSNN with approximation level 0.01 from 93% ( $\epsilon = 0$ ) to 71% (see labels E and F in Fig. 3) whereas the same attack drops the accuracy of the AccSNN from 96% to 82% (see Label G and H in Fig. 3). This indicates a 14% accuracy drop for the AccSNN, whereas the accuracy drop for the AxSNN is around 22% with PGD attack. In this case, the AxSNN is 1.5X more vulnerable than AccSNN.

Furthermore, we observe that the robustness of AxSNNs varies with the change in the approximation levels. For

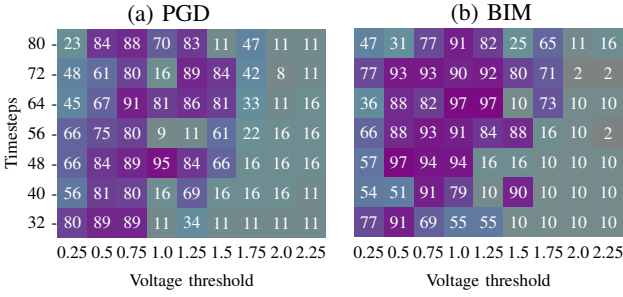


Figure 5: Accuracy of AxSNN (approx. level 0.01, precision-scale FP16) under attack ( $\epsilon = 1$ ) for MNIST.

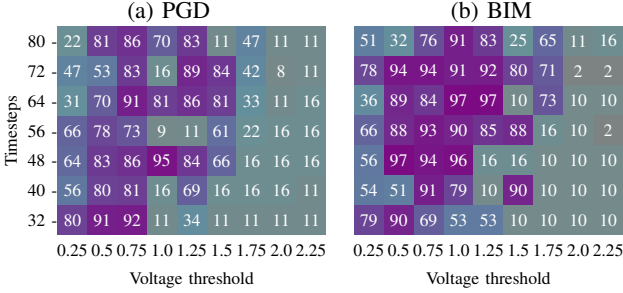


Figure 6: Accuracy of AxSNN (approx. level 0.01, precision-scale INT8) under attack ( $\epsilon = 1$ ) for MNIST.

instance, under no attack ( $\epsilon=0$ ), the accuracy of AxSNNs drops from 96% to 93%, 51%, and 10% when we change the approximate level from 0 to 0.01, 0.1 and 1.0, respectively. Conversely, the accuracy of AxSNNs drops from 96% ( $\epsilon=0$ ) to 77%, 25%, and 10% when we change the approximation levels from 0 to 0.01, 0.1 and 1.0, respectively under a PGD attack ( $\epsilon=0.9$ ) (see Fig. 2). A similar trend is also observed for the BIM attack (see Fig. 3).

Next, we explore the vulnerability of AxSNNs under sparse and frame attacks with the neuromorphic DVS128 Gesture dataset. We keep the threshold voltage and time steps constant to 1.0 and 80 in this experiment. Our results demonstrate that AxSNNs are vulnerable to neuromorphic attacks. For instance, the accuracy of AxSNN drops from 92% to 12% and 10% under sparse and frame attacks which is almost the same case in AccSNNs as shown in Fig. 7.

2) *Precision-scaling-based Adversarial Defense*: For evaluating the precision-scaling-based adversarial defense in AxSNNs for the static MNIST dataset, we start with the approximation level of 0.01 and vary the precision-scaling scales as Int8, FP16, and FP32. Specifically, we analyze the robustness of each precision-scaled AxSNN against PGD and BIM attacks, with perturbation budget ( $\epsilon=1.0$ ), by comparing the accuracy corresponding to each precision scale in Fig. 4, Fig. 5, and Fig. 6 with the base model accuracy (AccSNN) in Fig. 7a for the MNIST classification. We used  $\epsilon=1.0$  because the accuracy of both AccSNN and AxSNN drops significantly after this value, so it becomes non-recoverable. For example, it drops to 10% with  $\epsilon=1.5$ .

Furthermore, we observe that the robustness of precision-scaled AxSNNs varies while we change parameters such as threshold voltage, time steps, precision-scaling

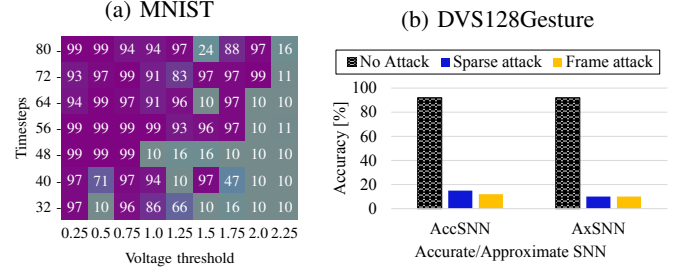


Figure 7: Accuracy of AccSNN without attack ( $\epsilon=0$ ) for MNIST and accuracy of AccSNN and AxSNN for DVS128 Gesture with and without attacks.

scale, and approximation level. For example, a PGD attack ( $\epsilon = 1.0$ ) results in a 12% accuracy loss in precision-scaled AxSNN with approximation level 0.01, precision scale FP32, threshold voltage 0.75 and time steps 32, when compared to AccSNN (see Fig. 7a). However, changing the precision scale to FP16 and INT8 results in 7% (see Fig. 5a) and 4% accuracy loss (see Fig. 6a, respectively, in the same AxSNNs when compared to AccSNN (see Fig. 7a). Likewise, changing threshold voltage and time steps to 1.0 and 48, respectively, results in recovering the accuracy of AxSNNs to 95%. Similarly, a BIM attack ( $\epsilon = 1.0$ ) results in a 15% accuracy loss in precision-scaled AxSNN with approximation level 0.01, precision-scaled FP32, threshold voltage 0.5 and time-step 72, when compared to AccSNN (see Fig. 7a). However, changing the precision scale to FP16 and INT8 results in 4% (see Fig. 5a) and 3% accuracy loss (see Fig. 6a, respectively, in the same AxSNN when compared to the AccSNN (see Fig. 7a). Similarly, changing threshold voltage and time steps to 1.25 and 64, respectively, results in recovering the accuracy of AxSNNs to 97%. This indeed highlights the importance of using our proposed Algorithm 1 for finding a robust sweet spot for the most suitable threshold voltage, time steps, approximation level, and precision scale under an attack. It is important to note that some deviating behavior from the above-discussed robustness trend is often observed at very high threshold voltage, i.e., greater than 1.75, and high approximation level, i.e., greater than 0.1. It is not surprising because SNNs typically lack performance efficiency for very high threshold voltage and high approximation levels since the spikes that trigger the LIF neurons may not cross the threshold voltage, resulting in a compromised performance of the whole network.

We test our proposed Algorithm 1 using threshold voltage in the range of 0.25 to 2.25 with an interval of 0.25, time steps in the range of 32 to 80 with an interval of 8, and precision scales as Int8, FP16, and FP32. As shown in Table I, our Algorithm 1 identifies the best parameter configurations even under PGD and BIM attacks. For example, our algorithm determines the approximation level 0.01 for precision-scaled AxSNN with the precision scale FP32, threshold voltage 1.0, and time steps 48 for accuracy of 97% even under PGD attack ( $\epsilon = 1.0$ ). However, it determines an approximation level 0.011 for precision-scaled AxSNN with the precision scale INT8, threshold voltage 0.75, and time steps 32 for accuracy



Table I: Best robustness settings for the proposed precision-scaled AxSNN classifier for the MNIST dataset.

$(V_{th}, T)$	Attacks	$(q_{th}, a_{th})$	Accuracy[%]
(0.25,32)	PGD	(FP32, 0.01)	88
	BIM	(INT8, 0.009)	80
(0.75,32)	PGD	(INT8, 0.011)	92
	BIM	(FP16, 0.013)	91
(1.0,48)	PGD	(FP32, 0.01)	<b>97</b>
	BIM	(INT8, 0.0125)	<b>96</b>

Table II: Recovered accuracy  $A_r$  and accuracy loss  $A_l$  when compared to baseline accuracy with AQF filtering in AxSNN  $(V_{th}, T) = (1.0, 80)$  using DVS128 Gesture.

Attacks	$(q_t, a_{th})$	$A_r$ [%]	$A_l$ [%]
Sparse Attack	(0.015, 0.1)	<b>90.01</b>	<b>2.0</b>
	(0.01, 0.15)	88.4	3.6
	(0.0, 0.001)	84.3	7.7
	(0.015, 0.1)	<b>91.1</b>	<b>1.0</b>
Frame Attack	(0.01, 0.15)	89.9	2.1
	(0.0, 0.001)	88.2	3.8

of 88% under the same attack. It is worth mentioning that the robustness of AxSNNs varies with the precision scales, and thus our algorithm provides different approximation levels for a combination of different threshold voltage, time steps, and precision scales. For instance, a PGD attack ( $\epsilon = 1.0$ ) on AxSNN with approximation level 0.1, time steps 32, and threshold voltage 0.25 results in a 72% accuracy loss, whereas the same attack on the precision-scaled AxSNN with approximation level 0.01, precision scale Int8 leads to only 17% accuracy loss in the static MNIST dataset (4X robustness improvement). Similar trend is observed with the BIM attack.

3) *AQF-based Adversarial Defense*: To evaluate the effectiveness of the proposed AQF-based adversarial defense in the precision-scaled AxSNN, we use Algorithm 2 for DVS128 Gesture classification under the sparse and frame attacks. Unfortunately, training AxSNNs takes a very long time; thus, we limit ourselves in testing the proposed Algorithm 2 with threshold voltage 1.0 and time steps 80 only. Note, these parameter settings are the most common in SNN research community for neuromorphic datasets [14]. Table II enlists the approximation levels identified by our algorithm for precision scales 0.015 and 0.01. We observe that a sparse attack on AxSNN leads to a 77% accuracy loss when compared to AccSNN without any attack. However, AQF-based adversarial defense in precision-scaled AxSNN recovers the accuracy close to the baseline accuracy. For instance, incorporating AQF-based adversarial defense with a precision scale of 0.015 and approximation level of 0.1 in a precision-scaled AxSNN under a sparse attack recovers the accuracy to almost 90%, which is only 2% less than the baseline accuracy. Interestingly, a frame attack on AxSNN leads to 82% accuracy loss when compared to AccSNN. However, AQF-based adversarial defense with a precision scale of 0.015 and an approximation level 0.1 in a precision-scaled AxSNN under a frame attack recovers the accuracy to almost 91%, which is only 1% less than the baseline accuracy. The reason behind such a tremendous improvement in robustness is that AQF masks noisy events that have a low correlation with each other in a perturbed neuromorphic dataset.

## VI. CONCLUSION

In this paper, we extensively explored the adversarial robustness of AxSNNs and proposed two novel defense methods: precision scaling and approximate quantization-aware filtering (AQF) for designing adversarially robust AxSNNs. To demonstrate the effectiveness of these defense methods, we employ the static MNIST and the neuromorphic DVS128 Gesture datasets. Our results show that AxSNNs are more prone to adversarial attacks than AccSNNs, but precision scaling and AQF significantly improve the robustness of AxSNNs. For instance, a PGD attack on AxSNN results in a 72% accuracy loss compared to AccSNN without any attack, whereas the same attack on the precision-scaled AxSNN leads to only a 17% accuracy loss in the static MNIST dataset (4X robustness improvement). Likewise, a Sparse Attack on AxSNN leads to a 77% accuracy loss when compared to AccSNN without any attack, whereas the same attack on an AxSNN with AQF leads to only a 2% accuracy loss in the neuromorphic DVS128 Gesture dataset (38X robustness improvement).

## VII. ACKNOWLEDGEMENTS

This work was partially supported by awards from U.S. Naval Research Lab under the grant N0017321C2016. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the U.S. Government or agency thereof.

## REFERENCES

- [1] A. Viale et al., "Carsnn: An efficient spiking neural network for event-based autonomous cars on the loihi neuromorphic research processor," in *IJCNN*. IEEE, 2021, pp. 1–10.
- [2] S. Sen et al., "Approximate computing for spiking neural networks," in *DATE*. IEEE, 2017, pp. 193–198.
- [3] A. Marchisio et al., "R-snn: An analysis and design methodology for robustifying spiking neural networks against adversarial attacks through noise filters for dynamic vision sensors," in *IROS*. IEEE, 2021, pp. 6315–6321.
- [4] A. Bhattacharjee et al., "Examining the robustness of spiking neural networks on non-ideal memristive crossbars," in *ISLPED*, 2022, pp. 1–6.
- [5] O. Nomura et al., "Robustness of spiking neural networks based on time-to-first-spike encoding against adversarial attacks," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 69, no. 9, pp. 3640–3644, 2022.
- [6] A. Marchisio et al., "Dvs-attacks: Adversarial attacks on dynamic vision sensors for spiking neural networks," in *IJCNN*. IEEE, 2021, pp. 1–9.
- [7] R. El-Allami et al., "Securing deep spiking neural networks against adversarial attacks through inherent structural parameters," in *DATE*. IEEE, 2021, pp. 774–779.
- [8] A. Siddique et al., "Is approximation universally defensive against adversarial attacks in deep neural networks?" in *DATE*. IEEE, 2022, pp. 364–369.
- [9] Y. LeCun et al., "Mnist handwritten digit database," *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, vol. 2, 2010.
- [10] A. Amir et al., "A low power, fully event-based gesture recognition system," in *CVPR*, 2017, pp. 7243–7252.
- [11] S. Venkataramani et al., "Axnn: Energy-efficient neuromorphic systems using approximate computing," in *ISLPED*. IEEE, 2014, pp. 27–32.
- [12] F. Khalid et al., "Qusecnet: Quantization-based defense mechanism for securing deep neural network against adversarial attacks," in *ISLPED*. IEEE, 2019, pp. 182–187.
- [13] R. V. W. Putra and M. A. Shafique, "Q-spinn: A framework for quantizing spiking neural networks," *IJCNN*, pp. 1–8, 2021.
- [14] E. P. Frady et al., "Neuromorphic nearest neighbor search using intel's pohoiki springs," in *NICE*, 2020, pp. 1–10.