

Sneaky Spikes: Uncovering Stealthy Backdoor Attacks in Spiking Neural Networks with Neuromorphic Data

Gorka Abad
 Radboud University, The Netherlands
 & Ikerlan Research Centre, Spain
 abad.gorka@ru.nl

Oğuzhan Ersoy
 Radboud University,
 The Netherlands
 oguzhan.ersoy@ru.nl

Stjepan Picek
 Radboud University,
 The Netherlands
 stjepan.picek@ru.nl

Aitor Urbieta
 Ikerlan Research Centre,
 Spain
 aurbietta@ikerlan.es

Abstract—Deep neural networks (DNNs) have demonstrated remarkable performance across various tasks, including image and speech recognition. However, maximizing the effectiveness of DNNs requires meticulous optimization of numerous hyperparameters and network parameters through training. Moreover, high-performance DNNs entail many parameters, which consume significant energy during training. To overcome these challenges, researchers have turned to spiking neural networks (SNNs), which offer enhanced energy efficiency and biologically plausible data processing capabilities, rendering them highly suitable for sensory data tasks, particularly in neuromorphic data. Despite their advantages, SNNs, like DNNs, are susceptible to various threats, including adversarial examples and backdoor attacks. Yet, the field of SNNs still needs to be explored in terms of understanding and countering these attacks. This paper delves into backdoor attacks in SNNs using neuromorphic datasets and diverse triggers. Specifically, we explore backdoor triggers within neuromorphic data that can manipulate their position and color, providing a broader scope of possibilities than conventional triggers in domains like images. We present various attack strategies, achieving an attack success rate of up to 100% while maintaining a negligible impact on clean accuracy. Furthermore, we assess these attacks’ stealthiness, revealing that our most potent attacks possess significant stealth capabilities. Lastly, we adapt several state-of-the-art defenses from the image domain, evaluating their efficacy on neuromorphic data and uncovering instances where they fall short, leading to compromised performance.

I. INTRODUCTION

Deep neural networks (DNNs) have achieved excellent performance in machine learning (ML) tasks in different domains, like computer vision [35], speech recognition [25], and text generation [7]. One key aspect of DNNs that has contributed to their success is their ability to learn from large amounts of data and discover complex patterns. This is achieved through multiple layers of interconnected neurons. The connections between these nodes are weighted and adjusted during training to minimize error and improve the model’s accuracy. DNNs have many hyperparameters that can be tuned to achieve top performance on a given task, but careful optimization of these

hyperparameters is crucial. Training a well-performing DNN can be time and energy expensive, requiring tuning many parameters with large training data. For example, training the GPT-3 model consumed about 190 000 kWh of electricity [12]. These models’ increasing complexity and computational requirements have led researchers to explore alternative approaches, such as spiking neural networks (SNNs) [22], [61], [15], [17].

SNNs can significantly reduce the energy consumption of DNNs. For instance, Kundu et al. [36] achieved better compute energy efficiency (up to 12.2×) compared to DNNs with a similar number of parameters. In addition to their energy efficiency, SNNs have several other benefits. SNNs can be more robust to noise and perturbations, making them more reliable in real-world situations [38]. More precisely, data obtained by a dynamic vision sensor (DVS) [58]—which SNNs can process—captures per pixel brightness changes asynchronously, instead of the absolute brightness in a constant rate—as in images. Compared to standard cameras, DVS cameras have low power consumption and capture low latency data, i.e., neuromorphic data, which also has high temporal resolution [10], [42]. Thus, SNNs can process data in a more biologically plausible manner, for example, by processing neuromorphic data, making them well-suited for tasks involving sensory data processing. In computer vision, significant advancements have been made in the context of autonomous driving, as evidenced by the exceptional performance attained [73]. The surrounding environment is captured by employing one or more vehicular cameras, with the following data being processed via a DNN. The decisions made by this DNN facilitate the autonomous operation of the vehicle. In recent works, alternative approaches have been proposed, wherein event-based neuromorphic vision is advocated to accomplish the same objective [10], [64]. Event-based data allows solving challenging scenarios where regular data (captured by standard cameras) cannot perform well [46], [76], such as high-speed scenes, low latency, and low power consumption scenarios. Moreover, SNNs are widely applicable, being used in domains like medical diagnosis [21] and computer vision [70], [27]. Finally, while DNNs are often considered to perform better than SNNs in terms of accuracy, recent results show this performance gap is reducing or even disappearing [61].

DNNs are vulnerable to various privacy and security threats, including adversarial examples [59], inference at-

tacks [8], model stealing [31], and backdoor attacks [26]. However, despite the widespread application domain of SNNs, their security aspects have yet to receive a comprehensive evaluation. Recent investigations [47] have conducted a comparative analysis of the security vulnerabilities inherent in SNNs to date, revealing their susceptibility to adversarial examples. Furthermore, subsequent research [63] has demonstrated that SNNs are also susceptible to hardware attacks, where the deliberate induction of bit-flips can lead to misclassification.

Backdoor attacks are a threat where malicious samples containing a trigger are included in the dataset at training time. After training, a backdoor model correctly performs the main task at test time while achieving misclassification when the input contains the trigger. Backdoor attacks on DNNs are well studied with several improvements such as stealthy triggers [45], [75] and dynamic triggers unique per data sample [49], [55]. Moreover, multiple works consider the backdoor defenses trying to detect and prevent backdoor attacks by inspecting DNN models [65], [43], [11], [40].

However, the existing backdoor attacks and defenses on DNNs do not directly apply to SNNs because of the different structures of SNNs and their usage of neuromorphic data. Unlike DNNs, SNNs do not have activation functions but spiking neurons, which could reduce or even disable the usage of existing attacks and defenses in DNNs that rely on them, as discussed in Section IV. Additionally, the time-encoded behavior of neuromorphic data allows a broader range of possibilities when generating input perturbations. At the same time, the captured data is encoded in much smaller pixel space (2-bit pixel space) than in regular images, which can handle up to 255 pixel possibilities per channel. The challenges regarding the application of backdoor attacks in SNNs are detailed in Section III.

To our knowledge, only one work explores backdoor attacks on SNNs [1]. The triggers used in the attack are static and moving square placed in the image during training time, which is not stealthy and is easily visible by human inspection, as investigated in Section V. Furthermore, their attack setup is limited, only considering three different poisoning rates and a single trigger size. Finally, the authors did not consider any backdoor defense.

This paper thoroughly investigates the viability of backdoor attacks in SNNs, the stealthiness of the trigger, and the robustness against the defenses. First, we improve the performance of the static and moving triggers proposed in [1]. Next, we propose two new trigger methods: smart and dynamic triggers. Our novel methods significantly outperform the existing backdoor attacks in SNNs. Our main contributions are:

- We explore different backdoor injecting methods on SNNs, achieving at least 99% accuracy in both main and backdoor tasks. We first explore static and moving triggers, which led to developing a smart attack that selects the optimal trigger location and color.
- We introduce the first dynamic backdoor attack on the neuromorphic dataset, which is highly stealthy.
- We analyze the stealthiness of backdoor attacks using the structural similarity index (SSIM) metric, showing that our dynamic trigger achieves up to 99.9% SSIM, outperforming static and moving triggers at 98.5% SSIM.

Additionally, we conduct a user study to measure the stealthiness of our attacks and compare their effectiveness with SSIM.

- We adapt image domain defenses for SNNs and neuromorphic data, observing their ineffectiveness against backdoor attacks.

We share our code to allow the reproducibility of our results.¹ Moreover, we show our triggers' dynamic motion and stealthiness as a live demo in the repository.

II. BACKGROUND

A. Backdoor Attacks

Backdoor attacks modify the behavior of a model during training, so at test time, it behaves abnormally [26]. A backdoored model misclassifies the inputs with a trigger while behaving normally on clean inputs. In a data poisoning backdoor, the training set is modified to include malicious data samples with the trigger. In the image domain, the trigger can be a pixel pattern in a specific image part. When the algorithm is trained on a mixture of clean and backdoor data, the model learns only to misclassify the inputs containing the pixel pattern, i.e., the trigger, to a particular target label.

Formally, an algorithm $f_{\theta}(\cdot)$ is trained on a mixture dataset containing clean and backdoor data, which rate is controlled by $\epsilon = \frac{m}{n}$ where n is the size of the clean dataset, m is the size of the backdoor dataset, and $m \ll n$. The backdoor dataset \mathcal{D}_{bk} is composed of backdoor samples $\{(\hat{\mathbf{x}}, \hat{y})\}^m \in \mathcal{D}_{bk}$, where $\hat{\mathbf{x}}$ is the sample containing the trigger and \hat{y} is the target label. For a clean dataset of size n , the training procedure aims to find θ by minimizing the loss function \mathcal{L} :

$$\theta' = \operatorname{argmin}_{\theta} \sum_{i=0}^n \mathcal{L}(f_{\theta}(\mathbf{x}_i), y_i), \quad (1)$$

where \mathbf{x} is input and y is label. During the training with backdoor data, Equation 1 is modified to include the backdoor behavior expressed as:

$$\theta' = \operatorname{argmin}_{\theta} \sum_{i=0}^n \mathcal{L}(f_{\theta}(\mathbf{x}_i), y_i) + \sum_{j=0}^m \mathcal{L}(f_{\theta}(\hat{\mathbf{x}}_j), \hat{y}_j).$$

B. Spiking Neural Networks & Neuromorphic Data

SNNs are a class of neural networks that model the activity of biological neurons by using discrete events, or spikes, to represent the output of a neuron when it reaches a certain threshold. In contrast to traditional artificial neural networks, which operate on continuous-valued signals, SNNs are event-driven and incorporate the temporal dynamics of neural activity [61]. SNNs comprise layers of interconnected spiking neurons, including input, hidden, and output layers—similar to DNNs.

Mathematically, the behavior of a spiking neuron can be modeled using the leaky integrate-and-fire (LIF) model, which describes the neuron's membrane potential as a function of time [30]. The membrane potential of a neuron is a result of the sum of its inputs. When the membrane potential reaches

¹<https://github.com/GorkaAbad/Sneaky-Spikes>

a certain threshold Θ , the neuron emits a spike, and its membrane potential is reset to a resting potential:

$$h(x) = \begin{cases} 1, & \text{if } x \geq \Theta \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

similar to the commonly used in deep learning (DL) rectified linear unit (ReLU) activation function [19]. However, note that Equation 2 is non-differentiable, which would lead to the impossibility of the derivative calculation during back-propagation, the *de facto* training algorithm in DNNs. Thus, there are two main options to train SNNs: spike timing-dependent plasticity (STDP) [2] or surrogate gradients [48]. STDP modifies the strength of connections between neurons based on the relative timing of their spikes. Specifically, if neuron A fires before neuron B , the connection between A and B is strengthened, whereas if neuron B fires before neuron A , the connection is weakened. This allows the network to adapt to changing input patterns and improve performance. Surrogate training, conversely, allows the approximation of the derivatives to perform backpropagation, enabling compatibility with Adam [33] or stochastic gradient descent [3]. The latter achieves better performance [38] and is the method we follow in our experimentation.

SNNs commonly operate on neuromorphic data, a time-encoded representation of the illumination changes of an object/subject captured by a DVS camera. The DVS camera captures a flow of spiking events that dynamically represents the changing visual scene. The advantage of using DVS cameras is that they provide a compressed representation of the visual scene that can be processed almost instantaneously, as reported in [58]. More precisely, neuromorphic data is encoded in T frames and p polarities. In neuromorphic data processing, polarity refers to the direction of the electrical signals representing data. These signals can be positive or negative, denoted as the *ON polarity* and *OFF polarity*, respectively; each polarity generates a different color. Using polarity in neuromorphic data processing can help reduce the energy required for computation. Using only positive or negative spikes reduces the number of bits required to represent data and thus reduces the system’s power consumption [61]. SNNs have shown promising results in various tasks [52], including speech recognition, image classification, and robotics.

III. BACKDOOR ATTACKS TO SNNs

A. Threat Model

We consider the same threat model as in prior studies [26], [13], [55], [1], [45], which assumes that the attacker can have full access to the model. Additionally, the attacker has access to the training dataset provided by the client. More precisely, we consider data poisoning-based, dirty label methodology for injecting the backdoor. We also limit our research to the digital image domain, where triggers are intended to be injected in digital samples rather than applied physically in the wild [5].

As a use case, we assume that a client wants to train an SNN on an owned dataset but does not have the resources, e.g., GPU cards, to train it. Therefore, the client outsources the training to a third party that provides on-cloud training

services, such as Google Cloud² or Amazon Web Services³, by sharing the model architecture and the training dataset. We assume that the attacker is the third-party provider, thus having access to the training procedure, the model, and the dataset. The attacker then injects the backdoor during training and shares the model with the client. The client can check the model’s performance using a holdout dataset.

B. Challenges in SNNs

SNNs have shown promising results in various domains, including image recognition [32] or image segmentation [51] with several applications ranging from autonomous driving to medical diagnosis, to name a few. However, as we demonstrate, SNNs are vulnerable to security threats, which can have serious consequences. SNNs have a unique network structure and utilize neuromorphic data. The information propagation is the key difference between “classical” DNNs and SNNs. SNNs do not work with continuously changing time values (like DNNs) but operate with discrete events that occur at certain points in time. The training is different, making the attacks happening in the training phase (potentially) challenging to deploy.

For instance, the triggers used in the image domain are encoded in 255 possibilities per channel, which gives many combinations of color. In neuromorphic data, however, the trigger space is reduced to 4 possibilities encoded by the two different polarities. Furthermore, in the image domain, the trigger is commonly static, i.e., no time-encoded data is used. In neuromorphic data, we encode the trigger using the time window, allowing us to create triggers that change the location through time. Next, we list the main challenges:

C.1: Designing and optimizing the trigger. Selecting and optimizing the trigger in the context of neuromorphic systems poses significant challenges due to the temporal nature and multiple frames per data point. The diverse range of options raises important questions: *How can we efficiently identify the trigger that maximizes the efficacy of the attack while minimizing its impact on clean accuracy? When designing a backdoor trigger specific to SNNs, what are the crucial parameters?*

C.2: Generating stealthy triggers. Generating stealthy triggers for neuromorphic systems presents difficulties since each trigger pixel can only assume four distinct values, making smoothing the trigger over clean data more complex than images with higher value ranges, such as 3×256 in regular data. This limitation raises the questions: *How can we design a backdoor trigger that exploits the time-encoded data to create a unique, imperceptible trigger for each input sample and frame? What influence do the selected parameters exercise on the stealthiness of the trigger?*

C.3: Backdoor defenses. Given that SNNs do not incorporate activation functions commonly employed in backdoor defense mechanisms, existing state-of-the-art defenses may not directly apply. These defenses are typically designed for image data, while neuromorphic data comprises multiple frames per data point and distinct color encoding. These dissimilarities lead to the following questions: *How effective are the current defense*

²<https://cloud.google.com>

³<https://aws.amazon.com>

strategies when applied to SNNs? How can we adapt these defenses to datasets encompassing multiple frames?

C.4: Assessing stealthiness. It is non-trivial to assess the stealthiness of a backdoor trigger via a subjective human perspective. Can we objectively assess the stealthiness of a trigger for neuromorphic data? If yes, how?

Knowing the limitations in color and the flexibility in changes through time, we propose different techniques for injecting a backdoor in SNNs. With this information, we improve two existing attacks (static and moving backdoors), and we propose two novel attacks: smart and dynamic backdoors.

C. Static Backdoor

Backdoor attacks in SNNs or neuromorphic data were not explored before the work of Abad et al. [1]. Inspired by backdoor attacks in the image domain [26], the authors replicated the square trigger used in neuromorphic data. By completely discarding the time-encoded advantages neuromorphic data have, the authors included the same trigger (same position and polarity) in all the frames, thus making a static trigger. We will investigate this trigger type as a baseline for subsequent ones, thoroughly investigating the trigger position, polarity, and size in a wide range of cases.

We follow the same intuition of a pixel-squared trigger of a given color, which is now set by the polarity for neuromorphic data. The data samples contain two polarity values, either *ON polarity* or *OFF polarity* corresponding to the black and light blue. However, when pixels from different polarities are overlapped, it generates another two color polarities, i.e., dark blue and green. The polarity p in neuromorphic datasets is a two-bit discrete value, creating up to four different combinations; we rename the polarities for simplicity to $p_0, p_1, p_2,$ and p_3 . Thus, the trigger gets a different color for different polarity combinations, i.e., black, dark blue, green, or light blue. Additionally, it can be placed in arbitrary locations of the input, e.g., top-right, middle-left, bottom-right, random, or any other desired location l , see Figure 1 as an example. For our attacks, we also consider the trigger size, s , as the percentage of the input size for constructing the trigger. Still, the input samples are divided into T frames, so the trigger k is replicated for each frame and sample, i.e., the trigger does not change the location. Consequently, it is static.

We consider all discussed parameters in the backdoor creation function $\mathcal{A}(\mathbf{x}, p, s, l)$ for creating a set of backdoor samples $\mathcal{D}_{bk} : \hat{\mathbf{x}} \in \mathcal{D}_{bk}$ containing the trigger k . By controlling the ϵ value $\epsilon = \frac{m}{n}$; $m \ll n$ with n the size of \mathcal{D}_{clean} and m the size of \mathcal{D}_{bk} , the attacker controls the amount of backdoored data during training.

D. Moving Backdoor

As previously seen, static backdoors replicate the trigger from backdoor triggers in the image domain. However, a unique characteristic of neuromorphic data allows the attacker to develop a better version of the trigger. To this end, moving triggers inject a trigger per frame in different locations, exploiting the time-encoded nature of neuromorphic data. The nature of neuromorphic data is driven by polarity, i.e., movement, which contradicts the static behavior of the naïve static attack. Driven by this discrepancy and the aim of creating




Figure 1: Input samples containing a static trigger (1a, 1b, and 1c) and a smart backdoor mask for $c = 2$ (1d).

a more stealthy attack that cannot be detected easily by human inspection (more about attack stealthiness is in Section V).

The moving backdoor primary leverages the “motion” nature of neuromorphic datasets to create moving triggers along the input. Precisely, for a given polarity p , a location l , and size s , the trigger k smoothly changes from frame to frame, creating a moving effect. Formally, the backdoor creation function takes the parameters mentioned above $\mathcal{A}(\mathbf{x}, p, s, l, T)$ for creating a backdoor set of inputs \mathcal{D}_{bk} , where T is the total number of frames that the input is divided. The backdoor creation function also considers the number of frames T , such that for each frame, $\mathcal{A}(\cdot)$ calculates a location at $t + 1 \in T$ close to the previous frame t . This allows the trigger to simulate a smooth movement in the input space. Unlike the static trigger, the moving trigger can be placed atop the input “activity area” for better stealthiness.

To create a moving trigger, we use a *binary mask*⁴. The value m has the same dimensions as the input \mathbf{x} , and it defines the trigger location and size. For example, given that at frame $t = 0$, the upper-left corner of the trigger is located in pixel (α, β) , the size of the trigger is s , the width of the image is w , and we shift the trigger 2 pixels to the right, then the mask is:

$$m(t)_{i,j} = \begin{cases} 1, & \text{if } i \in [\alpha \bmod w, (\alpha + s) \bmod w] \text{ and} \\ & j \in [(\beta + 2 \times t) \bmod w, (\beta + s + (2 \times t)) \bmod w] \\ 0, & \text{otherwise.} \end{cases}$$

We chose a shift of two pixels as it is used in the literature [1], and we verified through visual inspection that the trigger moved naturally in this way. If we used one pixel, the trigger moved very slowly; sometimes, it seemed static. Each pixel of the poisoned image $(\hat{\mathbf{x}}_{t,i,j})$ for every frame t is given by:

$$\hat{\mathbf{x}}_{t,i,j} = \begin{cases} \mathbf{x}_{t,i,j}, & \text{if } m_{t,i,j} = 0 \\ k_{i,j}, & \text{if } m_{t,i,j} = 1. \end{cases}$$

To improve previous work [1], we conducted a complete experimental setup to find the best moving trigger. Additionally, we analytically measure and quantify the stealthiness of the triggers. Finally, we correlate the stealthiness to the triggers’ ability to evade state-of-the-art defenses adapted from the image domain. Additional information about stealthiness can be found in Section V.

⁴A binary mask is a data structure, an image in this case, that consists of binary values (0 or 1) per pixel. A value of one means that a pixel from the trigger will be used, and when the value is zero, the input’s pixel is used. It is widely used in the literature to decide the position of the trigger [71], [14].

E. Smart Backdoor

So far, the proposed techniques, i.e., static and moving backdoors, inject the backdoor in the model correctly. With the *smart backdoor* approach, we aim to optimize the backdoor performance, simplicity, stealthiness, and understanding by removing the two hyperparameters: polarity p and trigger location l . For a better understanding of the effects of the trigger location, we split the image by drawing c vertical and horizontal lines (see Figure 1d). These will divide the image into $(c + 1)^2$ chunks, which we call *masks*; the collection of masks is denoted as \mathcal{S}_{mask} . Note that a larger c value would create more masks, allowing the attacker more control over the “optimal” spot of trigger placement. The smart backdoor attacks leverage the inputs’ polarity changes to find the most *active mask*. We define the *mask activation* as the sum of all the polarity changes happening in a mask for all the frames, excluding the polarity p_0 , which represents no movement, i.e., the black color or the background. This way, the smart backdoor finds the most active mask in the input sample. For example, the change from p_2 to p_1 or p_0 to p_1 is counted as 1. In contrast, the change from p_2 to p_0 does not count due to the representation of a “movement” state to a “no movement” state. Instead of calculating the activity per sample, we sum the activity over various samples. Note that per sample calculation of the activity would result in a sample-specific trigger, which is not the goal of the smart trigger. Sample-specific triggers are studied in Section III-F. We denote each polarity switch from p_n to p_m as $p_n \rightarrow p_m$. Formally, given a collection of masks, \mathcal{S}_{mask} , for each mask v_z where H is the height and W is the width of the mask:

$$f_{v_z} = \sum_{i=1}^W \sum_{j=1}^H \sum_{t=1}^T f_{i,j}$$

where $f_{i,j} = \begin{cases} 0 & \text{if } p_n \rightarrow p_0 \\ 1 & \text{otherwise,} \end{cases}$

where f_{v_z} demonstrates the calculated changes for the regions within that mask v_z during all frame steps t , and $f_{i,j}$ demonstrates the change just for pixel (i, j) in the image.

The most active mask v' is found by summing the activity of all the masks in the malicious set m :

$$v' = \operatorname{argmax}_{v_z} \sum_{j=1}^m f_{v_z}.$$

For example, in Figure 1d, $c = 2$ lines horizontally and vertically split the image into nine masks, and the smart attack will decide which mask is the most active. Let us take Figure 1d as an example, where the bottom middle mask, i.e., “2,c” is the most active one.⁵ Once the location is chosen, the smart attack also selects the best—performance-wise—polarity for the trigger, i.e., the least used polarity in the mask, denoted as p' . Therefore, p' is used for the trigger’s polarity k and is injected in v' , randomly and smoothly moving around the mask for all the frames.

Note that the trigger polarity p' and the mask v' are calculated for the poisoned dataset m , i.e., all the poisoned samples

⁵We also investigate the effect of trigger injection in the least active area in Section IV.

have the same trigger location and polarity. Formally, the smart backdoor creation function is defined as $\mathcal{A}(\mathbf{x}, p', v', T)$, generating a set of moving triggers that are combined with the clean samples to create a set of poisoned samples \mathcal{D}_{bk} .

Additionally, we investigate the effect of injecting the trigger in the least active masks, and we consider the usage of the most used polarity in the mask.

F. Dynamic Backdoor




Figure 2: Overview of the dynamic moving attack.

Having explored how to inject a backdoor in SNNs using static triggers, exploiting the time-encoded nature of neuromorphic data with moving backdoors, and optimizing the trigger polarity and location with the smart trigger, we propose a stealthy, invisible, and dynamic trigger. More precisely, motivated by dynamic backdoors in the image domain [49], [13], we investigate *dynamic moving backdoors* where the triggers are invisible, unique for each image and frame, i.e., the trigger changes from frame to frame. Note that generating a trigger that alternates in shape and color per sample and frame has not been previously investigated in the literature. Dynamic triggers are different from backdoors in the video domain as in [57], where the trigger is a patch, is not neuromorphic, does not change its location in time, and is not pursuing stealthiness. Neuromorphic data allows us to generate a dynamic trigger specific to a sample that is also unique per frame. To achieve this, we use a spiking autoencoder (AE)—which shares the common structure and usage as a standard autoencoder but with spiking neurons—to generate the optimal noise, as large as the image, which maximizes the backdoor performance, maintains a clean accuracy, and is invisible. More precisely, one of the weakest points of the previous backdoor triggers is that they are detectable under human inspection subject to the trigger location and polarity, see Section V. Therefore, we aim to create an invisible trigger that is not constrained by the polarity or the location. We leverage AE, which, for example, are used for denoising tasks, where we would usually require clean (denoised) and noisy versions of the image to train the AE, i.e., the AE is trained on image pairs [24]. However, we do not have the clean image and trigger pair to train the AE for our attack. If we had the trigger, there would be no need for the AE. Therefore, to fulfill these requirements, we must train the model and the AE simultaneously to make the AE generate a trigger unique for each sample and frame. Contrary to previous work [13], we do not need a fine-tuning phase to achieve a successful backdoor model, since attack success rate (ASR) and main task accuracy is already high. During training, we maximize the main task accuracy as well as the ASR, which is computationally more efficient.

Intuitively, the dynamic backdoor is designed as follows (see Figure 2). At first, we generate the perturbation by passing

a clean image to the spiking AE $g(\cdot) : \delta = g(\mathbf{x})$. The perturbation is then added to the clean image to construct a backdoor image $\hat{\mathbf{x}} = \mathbf{x} + \delta$. However, this naïve approach would saturate \mathbf{x} with δ , which makes the trigger visible. Thus, we project the perturbation to a l_p -ball⁶ of a given budget $\gamma : \|g(\mathbf{x})\|_\infty \leq \gamma$. Then, $g(\cdot)$ is updated aiming to maximize the backdoor accuracy of $f(\cdot)$, thus, during training, $g(\cdot)$ optimizes the parameters ζ that minimizes a loss function:

$$\zeta' = \underset{\zeta}{\operatorname{argmin}} \sum_{i=0}^n \mathcal{L}(f_\theta(g_\zeta(\mathbf{x}_i) + \mathbf{x}_i), \hat{y}_i),$$

$$s.t. \|g_\zeta(\mathbf{x})\|_\infty \leq \gamma \quad \forall \mathbf{x},$$

where \hat{y} is the target label, n is the length of the dataset, and \mathcal{L} is a loss function (mean squared error (MSE) in our case).

For training $f(\cdot)$, the parameters θ are updated by minimizing

$$\theta' = \underset{\theta}{\operatorname{argmin}} \sum_{i=0}^n \alpha \mathcal{L}(f_\theta(\mathbf{x}_i), y_i) +$$

$$(1 - \alpha) \mathcal{L}(f_\theta(g_\zeta(\mathbf{x}_i) + \mathbf{x}_i), \hat{y}_i),$$

$$s.t. \|g_\zeta(\mathbf{x})\|_\infty \leq \gamma \quad \forall \mathbf{x},$$

where α controls the trade-off between the clean and the backdoor performance; $\alpha = 1$ to train $f(\cdot)$ only with clean data. γ controls the visibility of the trigger. We discuss the influence of α and γ in Section IV.

IV. EVALUATION

A. Experimental Results

In this section, we provide the results for four different attacks, emphasizing their strong and weak points⁷. For details on the datasets, models, and training settings, we refer readers to Appendix, A-A, A-B, and A-C.

We evaluate the attacks with the commonly used metrics:

- **ASR** measures the backdoor performance of the model based on a holdout fully backdoored dataset.
- Model utility or **clean accuracy** is the performance of the model test on a holdout clean dataset.
- **Clean accuracy degradation** is the accuracy drop (in percentage) from the clean and backdoor models. It is calculated as $\frac{V_2 - V_1}{V_1} \times 100$, where V_1 is the clean baseline accuracy, and V_2 is the clean accuracy after the attack.

1) *Static Backdoor*: To first evaluate the viability of backdoor attacks in SNNs, we explore the basic BadNets [26] approach by placing a static trigger in different locations of the input space, using different poisoning rates, triggers sizes, and polarities. We test the static attack with ϵ values of 0.001, 0.005, 0.01, 0.05, and 0.1. We set the trigger sizes to 1% and 10% of the input image size. Lastly, we experiment with three trigger locations: bottom-right, middle, and top-left, and four polarities.

Our results show that static backdoors require a trigger size as big as 10% of the input size to inject the backdoor behavior in complex datasets like CIFAR10-DVS or DVS128-Gesture. When the trigger is 1% of the input size, the backdoor is only successfully injected in N-MNIST when the polarity is different from 0. However, when the trigger is in the middle, we observe that $p = 0$ achieves up to 100% ASR. This is caused because the data is centered; thus, the black trigger is on top of the image, contrasting and allowing the model to distinguish the trigger. In subsequent sections, we further investigate the importance of injecting the trigger in the most important or least important location. Increasing the trigger size makes the backdoor achieve an excellent ASR (up to 100%) when the trigger is placed in the corners. See Figure 10a, Figure 10b, and Figure 10c, in Appendix A-D for the results of top-left, middle, and bottom-right placed triggers.

Since the DVS128-Gesture dataset is small, the ϵ value drastically affects ASR. When $\epsilon = 0.01$, only a single sample will contain the trigger, which is insufficient to inject the backdoor when the trigger size is small. We further experiment with a larger trigger size, i.e., 0.3, achieving 99% ASRs with $\epsilon = 0.01$, in the top-left corner and using $p = 0$. CIFAR10-DVS achieves 100% ASR in all the settings when the trigger size and the poisoning rate are 0.1. CIFAR10-DVS is the only dataset that achieves 100% ASR in the bottom-right, with the polarity 0. This is caused by the dataset itself, which is noisy; thus, the black trigger can contrast with the background.

Regarding the clean accuracy degradation, we notice a slight degradation in most cases concerning the clean accuracy baseline. See Figure 12a, Figure 12b, and Figure 12c, for the results of the clean accuracy degradation of top-left, middle, and bottom-right placed triggers. N-MNIST does not show any degradation, i.e., 0% and in some cases even improving the accuracy by 0.1%, while DVS128-Gesture and CIFAR10-DVS are more prone to degrade the main task up to 5%. Overall, our static backdoors implementation in SNNs shows excellent performance, improving previous work [1] due to more precise tuning. Specifically, placing a static trigger in a moving input is unnatural, and it could be detected by checking if a part of the image is not moving or by inspecting changes in polarity between pixels. We address this limitation in the following sections.

2) *Moving Backdoor*: We investigate the effect of moving triggers to overcome the stationary behavior of static backdoor attacks. The moving backdoor changes the trigger position per frame horizontally, moving in a constant loop. We experiment with the same setting as static backdoors. However, the trigger location varies in time by horizontally moving using top-left, middle, and bottom-right as initial locations. The trigger changes location in two pixels every frame. Thus, the triggers change according to T—16 times in our experiments.

We observe that moving the backdoor overcomes the limitation of static triggers when placed on top of the image action. Since the trigger is moving, it is not always on top of the active area, thus allowing the model to capture both clean and backdoor features. Interestingly, as seen in Figure 11a and contrary to the static backdoor (see Figure 10a), triggers in the bottom-right corner with background polarity do not work with complex datasets because they merge better with

⁶ l_p -ball refers to a geometric shape that is defined as the set of all points in n -dimensional space that are within a certain distance of a given point, according to the l_p -norm.

⁷Due to space limitations, for more detailed results on the N-Caltech101 dataset, refer to Appendix A-K.

the image, making it impossible for the model to recognize them. However, in N-MNIST, we can achieve 100% ASRs with $p \neq 0$ and large ϵ . Moreover, triggers with background polarity in the bottom-right position do not inject the backdoor successfully for the DVS128-Gesture dataset, contrary to static backdoors. That effect is intuitively explained as moving backdoors are more difficult to inject than static ones. The model has to find a more complex relation between the trigger, samples, and label, thus requiring large datasets.

We observe the opposite behavior with triggers in the top-left corner (see Figure 11c). We investigate the data samples and conclude that images are usually centered; thus, the main action of the image is also in the middle of the image. For DVS128-Gesture or CIFAR10-DVS, the action is also contained in the corners of the image. From here, we can intuitively explain that injecting the trigger in an active or inactive area of the image could enable or disable the backdoor effect. We investigate the backdoor effect when placing the trigger in the most and least active areas in Section IV-A3. Lastly, by placing the triggers in the middle (see Figure 11b), we observe that the DVS128-Gesture dataset achieves 100% ASRs when polarity is 1 or 2. These results also suggest that the trigger’s polarity strongly affects the backdoor’s performance. Furthermore, depending on where the trigger is placed, a given polarity could have a different effect, as observed with polarity 2 in Figure 11b and Figure 11a for the DVS128-Gesture. Section IV-A3 investigates this effect in more detail.

Regarding the clean accuracy degradation, we notice a slight degradation in most cases concerning the clean accuracy baseline and even improving the accuracy of the main task in some settings. See Figure 13a, Figure 13b, and Figure 13c for the results of the clean accuracy degradation of top-left, middle, and bottom-right placed triggers. N-MNIST does not show any degradation, while DVS128-Gesture and CIFAR10-DVS are more prone to degrade the main task up to 7%, contrary to static triggers, which may degrade the clean accuracy up to 10%. We believe this happens as those datasets are more complex, and adding backdoors makes the main task more challenging.

3) *Smart Backdoor*: To explore the effects of the trigger location and the trigger polarity, we designed a novel attack that chooses the best combination of both. The smart attack removes the trigger location and the polarity selection by choosing either the most active or least active area of the image. Then, it chooses the least or most common polarity in that mask, where the least common polarity would contrast while the most common one would be more stealthy—enabling a more optimized attack. We experiment with different settings, such as poisoning rates, trigger sizes, and most and least active masks. We also investigate the trigger polarity’s effect using the most/least active polarity in the selected mask. We split the image using $c = 2$, see Figure 1d.

We observe that the backdoor is not successful with the DVS128-Gesture dataset. The few samples in the dataset make the choice of the most/least active mask of the image not precise. Note that the activity sum is done over all the images in the training set. The more samples, the more precise the selected mask is. Experimentation using the most active area shows excellent performance when the least common trigger polarity is used; see Figure 3a. Intuitively, the least common

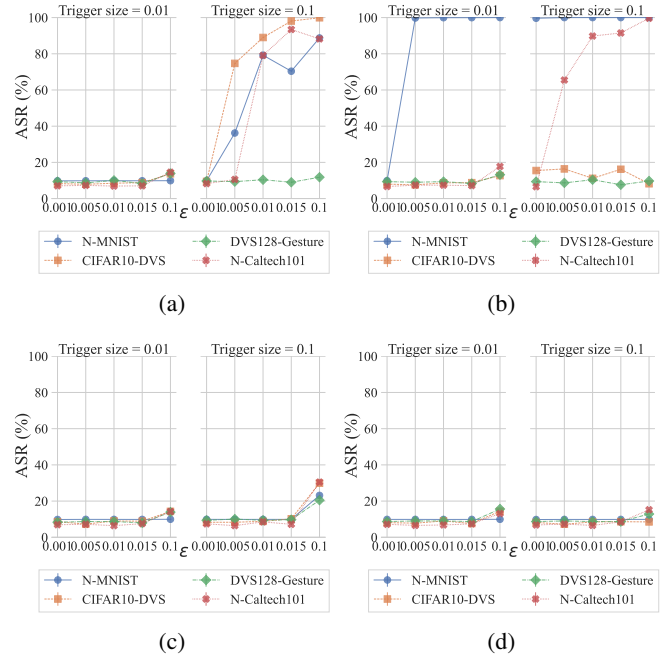


Figure 3: 3a and 3b show the smart triggers in the most active area, using the least and most common polarity. In 3c and 3d, we show the smart triggers in the least active area, using the least and most common polarity.

trigger polarity is preferred to increase ASR because of the trigger contrast compared to the background image.

Using 1% of the image size for the trigger only shows promising results with N-MNIST with $\epsilon = 0.1$. A larger trigger size improves the backdoor success using the least poisoned samples. Interestingly, injecting the trigger in the least active area with the least active trigger polarity, see Figure 3b, shows excellent backdoor performance even with a small trigger size. Finally, experimenting with the most active trigger polarities shows that the trigger merges with the actual image, not allowing the model to capture both the clean image and the trigger, see Figure 3c and Figure 3d.

Finding. The triggers are best injected in the most active area with the least common polarity.

We also experimented with the clean accuracy degradation with the smart trigger. As shown in Figure 14a, Figure 14b, Figure 15a, and Figure 15b, we observe a maximum of 4% degradation when the trigger size is 0.01 and the poisoning rate is 0.1 for the most active area. Results also show that when the trigger size is larger, the clean accuracy drop is negligible in all the cases, even improving it slightly. Injecting the trigger in the least active area shows similar performance; however, the accuracy degradation on the main task is smaller. This could be caused by the trigger not overlapping the active area, allowing the model to capture both tasks. This addresses Challenge C.1.

4) *Dynamic Backdoor*: In this study, we perform experiments using the dynamic backdoor technique to investigate the impact of different values of α and γ on the clean/backdoor trade-off and trigger intensity, respectively. We vary α in the




Figure 4: ASR and clean accuracy degradation of dynamic triggers. Dashed lines represent the ASR, and solid lines represent the clean accuracy. Blue corresponds to N-MNIST, orange to CIFAR10-DVS, green to DVS128-Gesture, and red to N-Caltech101.

range of 0.5 to 0.9 and γ in the range of 0.01 to 0.1, excluding $\alpha < 0.5$ due to its impracticality in real-world scenarios where assigning more weight to the backdoor than the clean task is unrealistic. We aim to maximize clean accuracy and achieve a high ASR by selecting epochs with higher values. It should be noted that the measurement accuracy across different epochs may introduce a larger standard error in the results.

The experimental results presented in Figure 4 demonstrate the effectiveness of our dynamic attack strategy. We achieve 100% ASR for all tested settings for the N-MNIST dataset while maintaining a clean, high accuracy. Particularly, when γ is small, no degradation in clean accuracy is observed. However, as γ increases, there is a noticeable decline in clean accuracy, as shown in Figure 4 and other datasets.

For the more complex DVS128-Gesture dataset, depicted in Figure 4, we achieve a minimum of 95% ASR (excluding $\alpha = 0.9$). However, we observe a slight degradation in clean accuracy when α is set to 0.5 or when γ is large. The reduction in ASR is most significant with $\alpha = 0.9$ and $\gamma = 0.01$ since it prioritizes the main task and triggers stealthiness. Moreover, as α increases, the drop in accuracy becomes less pronounced, reducing the impact of γ .

We observe lower performance in the case of CIFAR10-DVS (Figure 4). Here, γ plays a critical role in ASR and clean accuracy. The trigger became more visible with increasing γ , resulting in higher ASR and reduced clean accuracy. Conversely, we observe the opposite behavior when the trigger is nearly invisible ($\gamma = 0.01$). Similar to the DVS128-Gesture dataset, setting $\alpha = 0.9$ and $\gamma = 0.01$ reduces ASR. A clear trend is evident in this plot, where increasing α leads to a reduction in ASR while achieving higher clean accuracy. It is important to note that the baseline accuracy for CIFAR10-DVS is not as high as in the other datasets, which leads to more misclassification. This also incurs a more severe degradation of the clean accuracy after the attack.

Our findings indicate that ASR remains consistently close to 100% in most settings. However, the attacker must carefully select appropriate values for α and γ to balance clean accuracy

degradation and trigger invisibility. Notably, γ controls the visibility of the backdoor image, which becomes indistinguishable from the clean image when $\gamma = 0.01$. Our experiments suggest that using a small γ and a large α yield optimal results regarding clean accuracy and ASR while ensuring that the trigger remains unnoticed, thereby addressing Challenge C.2. Further details on the stealthiness aspect of our approach are discussed in Section V.

B. Evaluating State-of-the-Art Defenses

In this section, and due to the lack of specially crafted defenses for SNNs, we discuss state-of-the-art backdoor defenses in DNNs and how they could be adapted to SNNs and neuromorphic datasets. As discussed in the following sections, defenses for DNNs have core problems since they are based on DL assumptions or consider regular static data. We select four representative defenses based on model inspection: artificial brain stimulation (ABS) [43], STRIP [20], spectral signatures [62], and fine-pruning [41].

1) *ABS*: The ABS method, as introduced in Liu et al. [43], is a method for identifying backdoors in neural networks using a model-based approach. It works by stimulating neurons in a specific layer and examining the resulting outputs for deviations from expected behavior. ABS is based on the idea that a class can be represented as a subspace within a feature space and that a backdoored class will create a distinct subspace throughout the feature space. Therefore, ABS hypothesizes that a poisoned neuron activated with a target input will tend to produce a larger output than a non-poisoned neuron.

We adapt ABS to handle neuromorphic data and SNNs. Specifically, we modify the code to process all frames of an image together rather than treating each frame individually since neuromorphic data contains time-encoded information. However, ABS also does not support dynamic, moving, or smart triggers, which are types of backdoors that can change position or be unique to each image. Since the trigger position changes could be interpreted as multi-trigger backdoors—attacks that contain more than one trigger within a single input—ABS cannot handle them. Additionally, dynamic backdoors present a twofold problem for ABS. First, dynamic triggers can also be interpreted as multi-trigger. Second, ABS requires the trigger to be the same every time, i.e., the trigger is not unique per sample. However, the dynamic attack creates input-specific triggers, which can surpass ABS from its original design.

We observe several false positives when testing ABS against static backdoors. When applied to a clean model, ABS marked it as compromised, and when applied to a poisoned model, ABS identified it as compromised but with the wrong target class. This behavior was consistent across all datasets. One possible explanation for this issue is the core assumption of ABS. ABS relies on “turn-points” created by the activation functions in the model, such as ReLU. However, the lack of ReLU activation functions in SNNs makes ABS malfunction, providing inaccurate results.

2) *STRIP*: Unlike repairing or flagging a model as compromised, Gao et al. [20] investigated the detection of backdoor inputs during test time. The authors proposed a method called strong intentional perturbation (STRIP) to identify backdoor inputs at runtime by intentionally perturbing incoming inputs

and observing the randomness of predicted classes. Low entropy in the predicted classes indicates the presence of malicious input. The experiments conducted by the authors demonstrated that a decision boundary could effectively separate benign samples. The authors assumed access to a set of clean data, typically the test set, and created a set of backdoor data by interpolating different samples from the dataset. The entropy of clean and backdoor data was calculated, revealing their distinct separability. We adapted this mechanism to neuromorphic data. For constructing the poisoned test set, we performed frame-by-frame interpolation between samples, i.e., pairing the first frame of one sample with the first frame of another.

Our experiments demonstrated that the entropy levels of neuromorphic data are significantly lower than those of *regular* data, enabling a reasonably confident differentiation between clean and malicious samples. Figure 5 illustrates the entropy levels for various attacks and datasets, with entropy measured on each test set sample. For results on smart and moving triggers, refer to Figure 16 in the Appendix A-H. We derived three main observations from our experiments. First, the claim made by Gao et al. that poisoned data exhibits lower entropy than clean data does not always hold. In certain cases, clean data demonstrates lower entropy than poisoned data. Second, in the remaining cases, the entropy of clean and backdoor data overlaps, rendering them indistinguishable and inseparable. Third, overall entropy levels are much lower in neuromorphic data than in the regular data tested in [20]. For example, in the CIFAR-10 dataset, the mean entropy is approximately one, while in neuromorphic data, it is around 0.01.

3) *Spectral Signatures*: Recent research conducted by Tran et al. [62] has focused on mitigating dataset poisoning by identifying and eliminating compromised sub-populations within the dataset. The authors utilized statistical techniques such as singular value decomposition (SVD) to identify crucial input features and magnify the distribution difference in the latent space of the last convolutional layer. This approach facilitated the removal of the backdoor effect by retraining the network using clean data. It is important to note that this defense mechanism relies on having access to the compromised dataset, which may not be feasible in many scenarios.

We assessed the effectiveness of this mechanism against static and moving attacks.⁸ We selected attack parameters that achieved high accuracy in clean and backdoor tasks. Following the authors’ suggestion, we set the target label to 0 and the percentile to 85%. As shown in Table I, we observed no significant degradation or improvement in clean accuracy or ASR. However, we discovered that spectral signatures incorrectly flagged some legitimate samples as backdoors, resulting in compromised data remaining in the dataset used for model retraining. This aligns with our previous experience with STRIP, see Section IV-B2, where the entropy of clean and backdoor samples exhibited similarities. Further investigation is necessary to establish reliable mechanisms to effectively remove backdoor samples from compromised datasets.

4) *Fine-pruning*: Fine-pruning [41] is a defense mechanism against backdoor attacks composed of two parts: pruning

⁸Spectral signatures do not apply to dynamic backdoor attacks because the trigger is generated on the fly, and thus, the poisoned samples are not available for inspection.

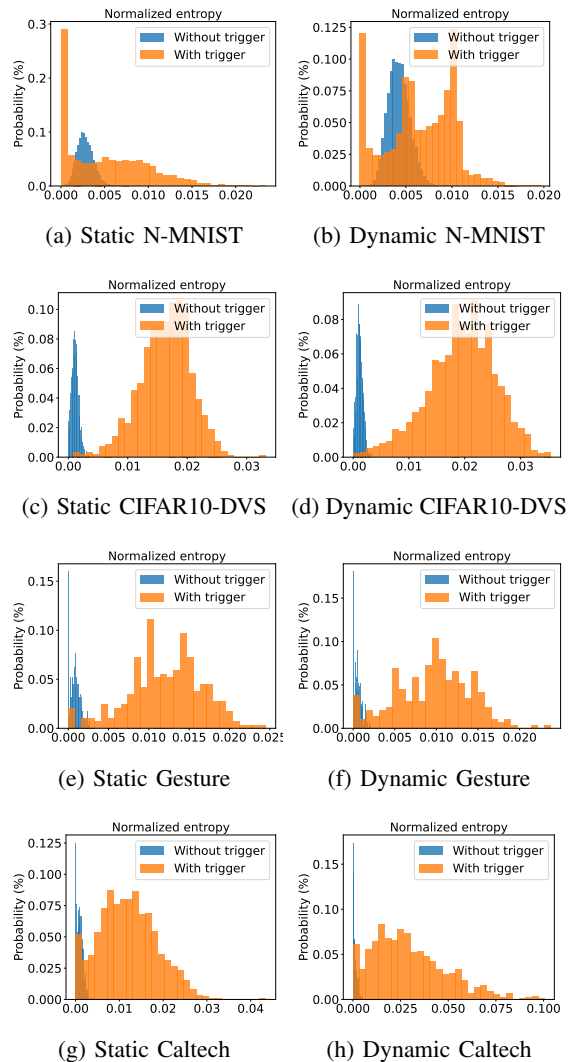


Figure 5: Normalized entropy of different triggers and datasets.

Table I: Comparison of the clean accuracy and the ASR between the baseline attack and after applying spectral signatures.

	Static				Moving			
	Baseline		Spectral		Baseline		Spectral	
	Clean acc.	ASR	Clean acc.	ASR	Clean acc.	ASR	Clean acc.	ASR
N-MNIST	99.4	100	99.4	100	99.3	100	99.3	100
CIFAR10-DVS	67.7	100	68.1	100	68.2	100	68.1	100
DVS128-Gesture	92.0	99.3	91.6	99.3	92.0	95.8	91.6	96.2
N-Caltech101	76.5	99.4	76.9	99.4	76.2	98.5	76.4	98.3

and fine-tuning. Existing works show that removing (pruning) some neurons of a DNN makes DNNs simpler, thus, easing the training while the prediction capacity remains equal [28], [72]. The authors suggested that some neurons may contain the primary task information, others the backdoor behavior, and the rest a combination of main and backdoor behavior. Thus, the backdoor could be completely removed by removing the neurons containing the malicious information. The authors proposed ranking the neurons in the last convolutional layer based on their activation values by querying some data. A pruning rate τ controls the number of neurons to prune.




Figure 6: Effect of fine-pruning on the ASR (dashed lines) and clean accuracy (full line) for different types of attacks, i.e., static, moving, smart, and dynamic. Blue corresponds to N-MNIST, orange to CIFAR10-DVS, green to DVS128-Gesture, and red to N-Caltech101.

The second part of the defense is fine-tuning. Fine-tuning consists of retraining the (pruned) model for some (small) number of epochs on clean data. By doing this, the model could (i) recover its dropped accuracy during pruning and (ii) altogether remove the backdoor effect. The authors showed that by combining these two, ASR of a poisoned model could drop from 99% to 0%.

We implemented this defense for SNNs and adapted it to work with neuromorphic data. We investigate the effect of pruning⁹, fine-pruning (pruning + fine-tuning), and fine-tuning only (when the pruning rate is 0). We also investigate various pruning rates, i.e., $\tau = \{0.0, 0.1, 0.3, 0.5, 0.8\}$ and analyze their impact, see Figure 6. Analyzing the results, we observe that pruning alone does not work. We notice that the clean accuracy drops drastically while ASR remains high, for example, as seen in Figure 17a. Depending on the trigger type, the drop in the clean accuracy is not that severe, but ASR remains high, as seen in Figure 17d. When combining pruning with a fine-tuning phase, i.e., fine-pruning, we observe that ASR can be drastically reduced while the clean accuracy remains high. The effect can be similar when focusing solely on the effect of fine-tuning, i.e., no pruned neurons ($\tau = 0$). Thus, pruning will not necessarily affect the model’s backdoor performance. However, we find that solely retraining the model with clean data reduces the backdoor effect. We conclude that fine-tuning could effectively reduce the backdoor performance while keeping clean accuracy high. Still, the effect is more pronounced for backdoors that aim to be more stealthy, making it an interesting trade-off. Since neuromorphic data consists of several frames (16), we can inject a moving backdoor, which is not possible in a single image. According to our experimental results shown in Figure 6b, we consider that fine-pruning may fail to reduce the effect of moving triggers.

Finding. Fine-pruning can be effective against backdoor attacks in SNNs using neuromorphic data. Still, this depends on the dataset characteristics and the trigger type.

The performance of the selected defenses in detecting backdoors in neuromorphic data and SNNs is limited by their inability to handle dynamic, moving, and smart triggers and their reliance on activation functions not present in SNNs. Further research and development are necessary to address these limitations and improve the robustness of the defenses in these contexts. Specific defenses considering the nature of

neuromorphic data and SNNs are necessary to detect and defend against these backdoors effectively. This addresses Challenge C.3.

a) Adaptive Attacker: Many defenses are intended to detect either malicious model parameters or samples by observing statistical differences between malicious and clean samples on (potentially) compromised models, e.g., neural cleanses (NC) [65] or fine-pruning [41]. By observing the effect of fine-pruning on our attacks, we investigated the ability of an adaptive attacker who knows the existence of defenses applied by the client in advance. Fine-pruning assumes that the backdoor effect is retained in some neurons while the clean behavior is retained in others. By stimulating the neurons in the last convolutional layer, the neurons with higher activation are thus compromised. The backdoor effect is removed by pruning the neurons with the highest activation (to some extent controlled by the pruning rate τ).

Recent work has made a substantial effort to develop techniques to bypass known defenses [60], [54]. We observed that fine-pruning results are ineffective by using a low poisoning rate. We experimented with $\epsilon = \{0.001, 0.01\}$ for all the datasets in different attack settings. In most studied cases (see Figure 7), the backdoor performance was kept high after fine-pruning, regardless of the pruning rate. Additionally, we investigate if the trigger size is relevant for fine-pruning. An adaptive attacker can bypass the defense even when increasing the trigger size to 30% of the image size and lowering the poisoning rate to 0.001: we test this with the DVS128-Gesture dataset. However, fine-pruning prevents the backdoor effect in the simplest case of a static trigger with N-MNIST $\epsilon = 0.01$ and a trigger size of 10%.

Following the same intuition of using a low poisoning rate in dynamic triggers, we can adjust the backdoor effect by tuning α . To bypass fine-pruning, we experiment with $\alpha = 0.9$ with the DVS128-Gesture and N-MNIST datasets as use cases. The results show similar behavior, and the attack maintains both clean high accuracy and ASR. Additionally, fine-pruning is often performed by pruning only in the last convolutional layer. An attacker who knows this beforehand could exclude this layer during training with backdoor data, so the rest of the layers contain the logic of the backdoor. Therefore, after pruning, the backdoor behavior will not be affected. Our experiments verify this hypothesis. A clean model trained on clean data on the DVS128-Gesture dataset achieved 90% accuracy. We retrain using backdoor data with $\epsilon = 0.1$, trigger size 10%, polarity 3, static trigger in the middle for 20 epochs, freezing the last convolutional layer and achieving 79% ASR. After fine-pruning, clean accuracy is 90%, and ASR is 89%. We also observe this behavior for CIFAR10-DVS with the same settings, achieving 68% clean accuracy and 100% ASR before fine-pruning and 67% clean accuracy and 100% ASR after fine-pruning.

V. STEALTHINESS EVALUATION

Quantifying image quality is often used for applications where the end user is a human. Subjective evaluation is commonly unsuitable for specific applications due to time constraints or expensive costs.

⁹For results on pruning, refer to Figure 17 in the Appendix.




Figure 7: Results of different adaptive attacks after fine-tuning. ASR (dashed lines) and clean accuracy (full line) for different types of attacks, i.e., static and moving. Blue corresponds to N-MNIST, orange to CIFAR10-DVS, green to DVS128-Gesture, and red to N-Caltech101.

A. User Study

We conducted a user study to validate the stealthiness of backdoor triggers identified using SSIM. The study aimed to assess the effectiveness of four triggers—static, moving, smart, and dynamic—on the DVS128-Gesture dataset. 25 participants—with different backgrounds in DL—including general researchers, practitioners, and students from various geographical locations, were recruited to participate in the user study.

During the study, participants were presented with a series of images, each containing one of the four backdoor triggers. Three images were clean, while one was compromised. To ensure a comprehensive evaluation, we explored various trigger positions and polarities for static and moving attacks, with the trigger occupying 10% of the image size, the largest in our experiments. We also assessed our smart attack at the least and most active locations and the least and most common polarities. Furthermore, our dynamic attack is evaluated for different γ values: $\gamma = 0.001, 0.01, 0.1$. The task assigned to the participants was to identify the compromised image.

The study recorded the selection frequency for each image as the compromised image by the participants. The analysis focused on calculating the percentage of times each image was chosen. The results revealed that the stealthiness of backdoor triggers varied based on location and color. The static, moving, and smart triggers exhibited varying stealthiness, heavily influenced by their parameter settings. An average of 50.6% of the participants correctly noticed the trigger in static settings. In moving backdoors, an average of 73.3% participants noticed the trigger correctly. An average of 84% of the participants found the trigger in smart triggers. As discussed in previous sections, we also found that the participants failed to select the correct poisoned sample when the trigger was placed in the image region with the most activity. The smart trigger’s stealthiness depends on the location and polarity. A smart trigger in the least common area with the least common polarity is more visible than a smart trigger in the most active location with the most common polarity. In contrast, the dynamic triggers displayed exceptional stealthiness, particularly at $\gamma = 0.01$, where only 4% of the participants found the trigger. At the same time, larger values of γ , like 0.1, render the triggers more visible, where 96% of the participants found the trigger.

These findings were consistent with the stealthiness evaluation using the SSIM metric.

The observed variations in stealthiness highlight the importance of considering the location, color, and parameter settings when assessing the effectiveness of backdoor triggers. The superior stealthiness exhibited by the dynamic triggers, which are not dependent on specific locations or colors, makes them particularly desirable in backdoor attack scenarios. Given that inspecting all data samples in a dataset is impractical, especially in real-life scenarios where models can be trained on billions of data samples [7], we evaluate the usability of metrics for quantifying the stealthiness of triggers in the subsequent sections.

Finding. A dynamic trigger cannot be detected by humans when generated using $\gamma = 0.01$.

B. Evaluated Metrics

MSE [66] compares two signals, e.g., image and audio, and measures the error or distortion between them. In our case, our signal is a frame sequence of images, where we compare a clean sample \mathbf{x} with a distorted (backdoored) sample $\hat{\mathbf{x}}$. In MSE, the error signal is given by $e = \mathbf{x} - \hat{\mathbf{x}}$, which is indeed the difference between pixels for two samples. However, MSE has no context neighbor pixels, which could lead to misleading results [23], [67]. For instance, a blurry image with an MSE score of 0.2, i.e., 20% of the pixels are modified, and a square of 20% of the sample size on top of the image would give the same MSE value. However, the blurry image is recognizable while the other is not. That is, two differently distorted images could have the same MSE for some perturbations more visible than others. Therefore, MSE cannot be the best measurement for backdoor attacks. Still, it could provide sufficient insights for quantifying stealthiness.

To overcome the locality of MSE, Wang et al. [68] proposed a measure called SSIM that compares local patterns of pixel intensities rather than single pixels, as in MSE. Images are highly structured, whereas pixels exhibit strong dependencies carrying meaningful information. SSIM computes the changes between two windows instead of the pixel-by-pixel calculations given by:

$$SSIM(\mathbf{x}, \hat{\mathbf{x}}) = \frac{(2\mu_{\mathbf{x}}\mu_{\hat{\mathbf{x}}} + c_1)(2\sigma_{\mathbf{x}\hat{\mathbf{x}}} + c_2)}{(\mu_{\mathbf{x}}^2 + \mu_{\hat{\mathbf{x}}}^2 + c_1)(\sigma_{\mathbf{x}}^2 + \sigma_{\hat{\mathbf{x}}}^2 + c_2)},$$

where $\mu_{\mathbf{x}}$ is the pixel sample mean of \mathbf{x} , $\mu_{\hat{\mathbf{x}}}$ is the pixel sample mean of $\hat{\mathbf{x}}$, and c_1 and c_2 are two variables to stabilize the division.

C. Evaluating Stealthiness

In this section, having analyzed metrics for comparing the variation between the clean and the backdoor images, we select SSIM as the most useful for our case. We evaluate the stealthiness of our different attacks based on the SSIM between the clean and backdoored images. The SSIM values are averaged over 16 (as the batch size) randomly selected images from the test set. Precisely, we compare each clean frame with its backdoor frame counterpart. Then, the SSIM per frame is averaged. To the best of our knowledge, this is the first application of SSIM for comparing similarities in neuromorphic data, used for backdoor attacks in SNNs, or used in the SNN domain overall. This addresses Challenge C.4.

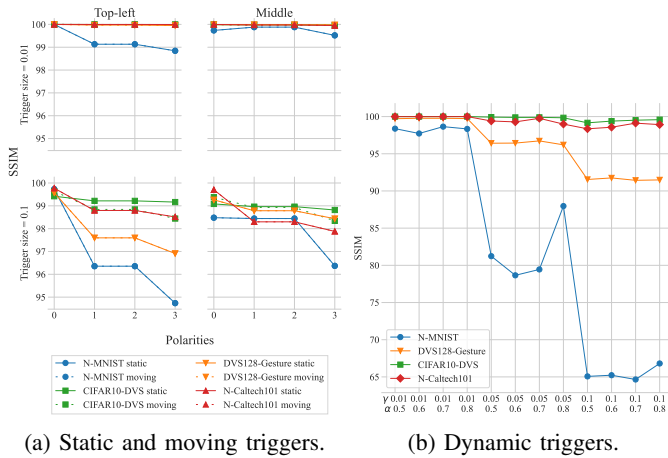


Figure 8: SSIM of different triggers.

1) *Static and Moving Triggers*: We first analyze the static and moving triggers in two positions: corner (top-left) and middle, see Figure 8a. We observe that the simpler the dataset, the more the stealthiness reduction, i.e., SSIM is lower. Additionally, the trigger size and the polarity affect the stealthiness. Indeed, the larger the trigger size, the less SSIM, which is expected. However, polarity also plays a crucial role in stealthiness. We observe a noticeable similarity downgrade related to the trigger polarity. The largest similarity degradation is observed for the N-MNIST dataset, with a trigger size of 0.01, and placing the trigger in the top-left corner. The background polarity, i.e., $p = 0$, shows high SSIM; however, with $p = 3$, the SSIM lowers to 94.5%. This could be directly linked to the number of pixels of a given polarity in an area. The less polarity in an area, the more “contrast” it would create, being less stealthy.

Comparing the static and moving triggers, we observe a more significant degradation when the trigger is moving, although it is negligible in some datasets or settings. Triggers in noisy datasets like CIFAR10-DVS and N-Caltech101 are more tolerable to input perturbations. Modifications in simpler datasets, such as N-MNIST, significantly change the image’s overall structure, achieving a lower SSIM.

2) *Smart Triggers*: Smart triggers select the trigger polarity and location by themselves, based on the image’s least or most active area and the least prominent polarity in an area. We observe a larger degradation based on the trigger size and when placed in the least active area (see Figure 18 in Appendix A-J). This is expected as the trigger in the most active area gets hidden by the high activity, i.e., motion. Thus, the performance of triggers in the most active area gets lowered, but gains trigger stealthiness, which must be considered a trade-off between stealthiness and performance.

3) *Dynamic Triggers*: Lastly, dynamic triggers (see Figure 8b) show impressive stealthiness as γ gets smaller, even to the point of being indistinguishable from the clean image. We also observe that the more complex the dataset, the less the reduction in the stealthiness, contrary to N-MNIST, where the degradation is notable with $\gamma = 0.1$. This effect is related to the number of pixel changes and the noise in the data.

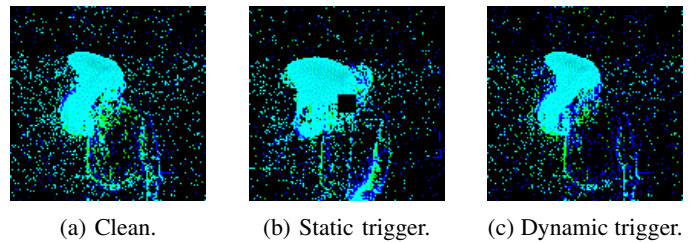


Figure 9: Comparison of triggers.

A dataset with large noise has much activity, thus making it easier for the trigger to be hidden, as in CIFAR10-DVS and N-Caltech101. However, with “clean” datasets that contain little noise as N-MNIST, even the subtlest perturbation makes a noticeable change. Still, with $\gamma = 0.01$, the perturbation in every tested dataset is invisible.

Overall, note that even if the SSIM of the static trigger (Figure 8a) and the dynamic trigger (Figure 8b) are similar, as seen in Figure 9 the visibility of the triggers are rather different. The static trigger is highly noticeable in the middle of the figure, while the dynamic trigger is indistinguishable from the clean sample. Although SSIM has been previously used for comparing images, we require more robust metrics to compare neuromorphic data.

In general, employing metrics to identify outliers in data similarity holds potential for practical applications and large-scale scenarios. However, it should be noted that while certain evaluations using SSIM yield comparable outcomes to those observed in our user study, SSIM alone cannot serve as an ad-hoc defense against backdoor attacks. This approach tends to produce many false negatives and can be easily circumvented by sophisticated adversaries. Addressing this limitation is an essential direction for our future research.

VI. RELATED WORK

SNNs. In recent years, numerous efforts have been made to develop supervised learning algorithms for SNNs to make them more practical and widely applicable. One of the first such algorithms was Spike Prop [6], which was based on backpropagation and could be used to train single-layer SNNs. However, it was not until more recent developments that SNNs could be applied to multi-layer setups. Despite such advances, most existing SNN training methods still require manual tuning of the spiking neuron membrane, which can be time-consuming and may limit the performance of SNN. To overcome this limitation, Fang et al. [17] proposed a method that can learn the weights and hyperparameters of the membranes in an automated way, eliminating the need for manual tuning. This advancement may make SNNs more practical and easier to use for a broader range of applications.

Several other notable developments in the field of SNNs are worth mentioning. One such development is event-driven update rules, allowing SNNs to operate more efficiently by only updating the network when necessary [74]. This contrasts traditional neural networks that require continuous updates and can be computationally expensive. Another area of research in

SNNs is structural plasticity, which refers to the network’s ability to change its structure during training [74], [69]. This can be accomplished through the addition or removal of connections between neurons or through the creation of new neurons altogether. Structural plasticity can improve the learning efficiency and generalization capabilities of SNNs and is effective in various tasks. There are also ongoing efforts to develop unsupervised learning algorithms for SNNs, allowing them to learn from data without needing labeled examples [29]. Unsupervised learning is a critical component of the brain’s learning process and can significantly expand the range of tasks that SNNs can perform.

Backdoor Attacks. Backdoor attacks were first introduced by Gu et al., where the authors presented BadNets [26]. BadNets uses a square-shaped trigger on a fixed location to inject the backdoor task; it was the first to show backdoor vulnerabilities in machine learning. BadNets requires access to the training data for injecting the backdoor, contrary to the work by Liu et al. [44], which alleviated this assumption. The authors presented a novel work where access to the training data was not needed. The authors systematically reconstructed training samples to inject the backdoor by adding the trigger on top of the samples and retraining the model. The discussed approaches use static triggers, i.e., the trigger is in the exact location for all the samples. Nguyen and Tran [49] developed a dynamic backdoor attack in which the trigger varies with the input. Specifically, a generator creates a pixel scatter that is then overlapped with the clean input. A similar approach was investigated by Salem et al. [55], who also constructed a dynamic backdoor attack. Instead of a pixel-scattered trigger, the trigger is a square, which has the advantage of applying it to physical objects. Aiming to increase the stealthiness of the backdoor, Lie et al. created ReFool [45], which includes benign-looking triggers as reflections in the clean sample. A similar approach was followed by Zhang et al. [75], who proposed Poison Ink, where the structure of the image is extracted. The structure is then injected with a crafted pixel pattern and included in the clean image. The resulting poisoned image is indistinguishable from the clean sample.

In SNNs and neuromorphic datasets, only [1] explored backdoor attacks. However, their experimentation is limited to exploring static and moving triggers using simple datasets and models. Moreover, based on the results, the authors do not provide insights into why backdoors occur in SNNs. They also do not consider any more advanced triggers, stealthiness evaluation, or defense mechanisms.

Defenses Against Backdoor Attacks. Backdoor attacks can be mitigated using either model-based or data-based defenses. Model-based defenses involve examining potentially infected models to identify and reconstruct the backdoor trigger. An example is NC [65], which aims to reconstruct the smallest trigger capable of causing the model to misclassify a specific input. This approach is based on the premise that an infected model is more likely to misclassify an input with a trigger than one that does not have a trigger. Another model-based approach is ABS [43], which involves activating every model neuron and analyzing anomalies in the resulting output. Model-based defenses are designed to specifically target infected models, searching for signs of a trigger and attempting to reconstruct it. This approach is effective because

the presence of a trigger is often a reliable indicator that the model has been compromised. However, it is also possible for a model to be infected without a trigger, where they are not applicable. Thus, model-based defenses are limited to cases with triggers. Data-based defenses aim to detect the presence of a backdoor by analyzing the dataset without inspecting the model itself. One approach in this category is clustering techniques to differentiate between clean and poisoned data [9]. Another approach, STRIP [20], combines the provided dataset with known infected data and queries the model on the resulting combined dataset. By measuring the entropy of the model’s output on this combined dataset, STRIP can detect backdoored inputs, which tend to have lower entropy than clean inputs. Data-based defenses focus on analyzing the dataset to detect the presence of a backdoor. These approaches are helpful because they do not require knowledge of the specific trigger used to infect the model, making them more robust against variations in the method of infection. However, data-based defenses may not be as effective at detecting more subtle forms of backdoor attacks, which may not leave as clear a signature in the dataset. Currently, no defenses are SNN or neuromorphic data-specific. As shown, some well-performing defenses adapted from the image domain do not work well in SNNs. Thus, developing SNNs or neuromorphic data-specific defenses is necessary for future research.

VII. CONCLUSIONS & FUTURE WORK

This study explored the security of SNNs in the context of backdoor attacks. Despite the growing importance of SNNs as an emerging technology, this area has received limited attention. Our investigation utilizes neuromorphic data and triggers to launch backdoor attacks in SNNs. We proposed several attack methods, including a novel dynamic trigger that evolves over time and remains undetectable to human inspection. We have also evaluated the attacks against different state-of-the-art defenses, which we have adapted from the image domain. Our results demonstrate that our attacks can achieve an ASR of up to 100% without causing noticeable degradation in clean accuracy, even when the defenses are employed. Our findings show that SNNs are highly vulnerable to backdoor attacks, indicating a need for further improvements in existing defense mechanisms.

Future research should focus on developing SNNs specifically designed to counter backdoor attacks and address the unique challenges posed by neuromorphic data. It is worth noting that SNNs are also gaining prominence in other domains, such as the graph domain, where backdoor attacks are becoming increasingly significant. In these domains, specific backdoor designs may be required to tackle the challenges that the spiking graph neural networks pose. Finally, while our study has primarily focused on backdoor attacks, it is important to consider other common threats that can be adapted to SNNs, such as inference attacks and adversarial examples.

ACKNOWLEDGMENT

We thank the shepherd for helpful feedback and support in improving this work. The Horizon Europe, Spanish CDTI, and ELKARTEK programs fund this research. Grant agreements 101021911 (IDUNN), CER-20191012 (EGIDA), and KK-2021/00091 (REMEDY), respectively.

REFERENCES

- [1] G. Abad, O. Ersoy, S. Picek, V. J. Ramírez-Durán, and A. Urbieto, "Poster: Backdoor attacks on spiking nns and neuromorphic datasets," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 3315–3317.
- [2] L. F. Abbott and S. B. Nelson, "Synaptic plasticity: taming the beast," *Nature Neuroscience*, vol. 3, no. 11, pp. 1178–1183, Nov. 2000. [Online]. Available: https://www.nature.com/articles/nn1100_1178
- [3] S.-i. Amari, "Backpropagation and stochastic gradient descent method," *Neurocomputing*, vol. 5, no. 4-5, pp. 185–196, 1993.
- [4] A. Amir, B. Taba, D. Berg, T. Melano, J. McKinstry, C. Di Nolfo, T. Nayak, A. Andreopoulos, G. Garreau, M. Mendoza *et al.*, "A low power, fully event-based gesture recognition system," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7243–7252.
- [5] E. Bagdasaryan and V. Shmatikov, "Blind backdoors in deep learning models," in *Usenix Security*, 2021.
- [6] S. M. Bohte, J. N. Kok, and H. La Poutre, "Error-backpropagation in temporally encoded networks of spiking neurons," *Neurocomputing*, vol. 48, no. 1-4, pp. 17–37, 2002.
- [7] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [8] N. Carlini, F. Tramèr, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, U. Erlingsson *et al.*, "Extracting training data from large language models," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 2633–2650.
- [9] B. Chen, W. Carvalho, N. Baracaldo, H. Ludwig, B. Edwards, T. Lee, I. Molloy, and B. Srivastava, "Detecting backdoor attacks on deep neural networks by activation clustering," *arXiv preprint arXiv:1811.03728*, 2018.
- [10] G. Chen, H. Cao, J. Conrath, H. Tang, F. Rohrborn, and A. Knoll, "Event-based neuromorphic vision for autonomous driving: A paradigm shift for bio-inspired visual sensing and perception," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 34–49, 2020.
- [11] Z. Chen, S. Wang, A. Fu, Y. Gao, S. Yu, and R. H. Deng, "Linkbreaker: Breaking the backdoor-trigger link in dnns via neurons consistency check." IEEE, 2000.
- [12] P. Dhar, "The carbon impact of artificial intelligence." *Nat. Mach. Intell.*, vol. 2, no. 8, pp. 423–425, 2020.
- [13] K. Doan, Y. Lao, W. Zhao, and P. Li, "Lira: Learnable, imperceptible and robust backdoor attacks," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 11 966–11 976.
- [14] Y. Dong, X. Yang, Z. Deng, T. Pang, Z. Xiao, H. Su, and J. Zhu, "Black-box detection of backdoor attacks with limited information and data," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 16 482–16 491.
- [15] J. K. Eshraghian, M. Ward, E. Neftci, X. Wang, G. Lenz, G. Dwivedi, M. Bennamoun, D. S. Jeong, and W. D. Lu, "Training spiking neural networks using lessons from deep learning," *arXiv preprint arXiv:2109.12894*, 2021.
- [16] W. Fang, Y. Chen, J. Ding, D. Chen, Z. Yu, H. Zhou, Y. Tian, and other contributors, "Spikingjelly," <https://github.com/fangwei123456/spikingjelly>, 2020, accessed: 2022-10-12.
- [17] W. Fang, Z. Yu, Y. Chen, T. Masquelier, T. Huang, and Y. Tian, "Incorporating learnable membrane time constant to enhance learning of spiking neural networks," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 2661–2671.
- [18] L. Fei-Fei, R. Fergus, and P. Perona, "Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories," in *2004 conference on computer vision and pattern recognition workshop*. IEEE, 2004, pp. 178–178.
- [19] K. Fukushima, "Cognitron: A self-organizing multilayered neural network," *Biological cybernetics*, vol. 20, no. 3-4, pp. 121–136, 1975.
- [20] Y. Gao, C. Xu, D. Wang, S. Chen, D. C. Ranasinghe, and S. Nepal, "Strip: A defence against trojan attacks on deep neural networks," in *Proceedings of the 35th Annual Computer Security Applications Conference*, 2019, pp. 113–125.
- [21] S. Ghosh-Dastidar and H. Adeli, "Improved spiking neural networks for eeg classification and epilepsy and seizure detection," *Integr. Comput.-Aided Eng.*, vol. 14, no. 3, p. 187–212, aug 2007.
- [22] S. Ghosh Dastidar and H. Adeli, "Spiking neural networks," *International journal of neural systems*, vol. 19, no. 04, pp. 295–308, 2009.
- [23] B. Girod, "What's wrong with mean-squared error?" *Digital images and human vision*, pp. 207–220, 1993.
- [24] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [25] A. Graves, A.-r. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *2013 IEEE international conference on acoustics, speech and signal processing*. Ieee, 2013, pp. 6645–6649.
- [26] T. Gu, K. Liu, B. Dolan-Gavitt, and S. Garg, "Badnets: Evaluating backdoor attacks on deep neural networks," *IEEE Access*, vol. 7, pp. 47 230–47 244, 2019.
- [27] A. Gupta and L. N. Long, "Character recognition using spiking neural networks," in *2007 International Joint Conference on Neural Networks*, 2007, pp. 53–58.
- [28] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding," in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2016. [Online]. Available: <http://arxiv.org/abs/1510.00149>
- [29] H. Hazan, D. Saunders, D. T. Sanghavi, H. Siegelmann, and R. Kozma, "Unsupervised learning with self-organizing spiking neural networks," in *2018 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018, pp. 1–6.
- [30] E. Hunsberger and C. Eliasmith, "Spiking Deep Networks with LIF Neurons," Tech. Rep., Oct. 2015, arXiv:1510.08829 [cs] type: article. [Online]. Available: <http://arxiv.org/abs/1510.08829>
- [31] M. Jagielski, N. Carlini, D. Berthelot, A. Kurakin, and N. Papernot, "High accuracy and high fidelity extraction of neural networks," in *29th USENIX security symposium (USENIX Security 20)*, 2020, pp. 1345–1362.
- [32] S. R. Kheradpisheh, M. Ganjtabesh, S. J. Thorpe, and T. Masquelier, "STDP-based spiking deep convolutional neural networks for object recognition," *Neural Networks*, vol. 99, pp. 56–67, Mar. 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0893608017302903>
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [34] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [35] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, 2012.
- [36] S. Kundu, G. Datta, M. Pedram, and P. A. Beerel, "Spike-thrift: Towards energy-efficient deep spiking neural networks by limiting spiking activity via attention-guided compression," in *Proceedings of the IEEE/CVF WACV*, 2021, pp. 3953–3962.
- [37] Y. LeCun, "The mnist database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998.
- [38] J. H. Lee, T. Delbruck, and M. Pfeiffer, "Training deep spiking neural networks using backpropagation," *Frontiers in neuroscience*, vol. 10, p. 508, 2016.
- [39] H. Li, H. Liu, X. Ji, G. Li, and L. Shi, "Cifar10-dvs: an event-stream dataset for object classification," *Frontiers in neuroscience*, vol. 11, p. 309, 2017.
- [40] Y. Li, X. Lyu, N. Koren, L. Lyu, B. Li, and X. Ma, "Neural attention distillation: Erasing backdoor triggers from deep neural networks," 2021.
- [41] K. Liu, B. Dolan-Gavitt, and S. Garg, "Fine-pruning: Defending against backdoor attacks on deep neural networks," in *International Symposium on Research in Attacks, Intrusions, and Defenses*. Springer, 2018, pp. 273–294.
- [42] S.-C. Liu, B. Rueckauer, E. Ceolini, A. Huber, and T. Delbruck, "Event-driven sensing for efficient perception: Vision and audition algorithms," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 29–37, 2019.

- [43] Y. Liu, W.-C. Lee, G. Tao, S. Ma, Y. Aafer, and X. Zhang, "Abs: Scanning neural networks for back-doors by artificial brain stimulation," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 1265–1282.
- [44] Y. Liu, S. Ma, Y. Aafer, W.-C. Lee, J. Zhai, W. Wang, and X. Zhang, "Trojaning attack on neural networks," 2018.
- [45] Y. Liu, X. Ma, J. Bailey, and F. Lu, "Reflection backdoor: A natural backdoor attack on deep neural networks," in *European Conference on Computer Vision*. Springer, 2020, pp. 182–199.
- [46] A. I. Maqueda, A. Loquercio, G. Gallego, N. García, and D. Scaramuzza, "Event-based vision meets deep learning on steering prediction for self-driving cars," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 5419–5427.
- [47] A. Marchisio, G. Nanfa, F. Khalid, M. A. Hanif, M. Martina, and M. Shafique, "Is Spiking Secure? A Comparative Study on the Security Vulnerabilities of Spiking and Deep Neural Networks," in *2020 International Joint Conference on Neural Networks (IJCNN)*, Jul. 2020, pp. 1–8, iSSN: 2161-4407.
- [48] E. O. Neftci, H. Mostafa, and F. Zenke, "Surrogate Gradient Learning in Spiking Neural Networks: Bringing the Power of Gradient-Based Optimization to Spiking Neural Networks," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 51–63, Nov. 2019.
- [49] T. A. Nguyen and A. Tran, "Input-aware dynamic backdoor attack," *Advances in Neural Information Processing Systems*, vol. 33, pp. 3454–3464, 2020.
- [50] G. Orchard, A. Jayawant, G. K. Cohen, and N. Thakor, "Converting static image datasets to spiking neuromorphic datasets using saccades," *Frontiers in neuroscience*, vol. 9, p. 437, 2015.
- [51] K. Patel, E. Hunsberger, S. Batir, and C. Eliasmith, "A Spiking Neural Network for Image Segmentation," Tech. Rep., Jun. 2021, arXiv:2106.08921 [cs] type: article. [Online]. Available: <http://arxiv.org/abs/2106.08921>
- [52] F. Ponulak and A. Kasinski, "Introduction to spiking neural networks: Information processing, learning and applications," *Acta neurobiologiae experimentalis*, vol. 71, no. 4, pp. 409–433, 2011.
- [53] C. Posch, D. Matolin, and R. Wohlgenannt, "High-dr frame-free pwm imaging with asynchronous aer intensity encoding and focal-plane temporal redundancy suppression," in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. IEEE, 2010, pp. 2430–2433.
- [54] X. Qi, T. Xie, Y. Li, S. Mahloujifar, and P. Mittal, "Revisiting the Assumption of Latent Separability for Backdoor Defenses," Feb. 2023. [Online]. Available: https://openreview.net/forum?id=_wSHsgrVali
- [55] A. Salem, R. Wen, M. Backes, S. Ma, and Y. Zhang, "Dynamic backdoor attacks against machine learning models," in *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2022, pp. 703–718.
- [56] A. Samadzadeh, F. S. T. Far, A. Javadi, A. Nickabadi, and M. H. Chehreghani, "Convolutional spiking neural networks for spatio-temporal feature extraction," *arXiv preprint arXiv:2003.12346*, 2020.
- [57] T. Sato, J. Shen, N. Wang, Y. Jia, X. Lin, and Q. A. Chen, "Dirty road can attack: Security of deep learning based automated lane centering under {Physical-World} attack," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 3309–3326.
- [58] T. Serrano-Gotarredona and B. Linares-Barranco, "A 128×128 1.5% contrast sensitivity 0.9% fpn $3 \mu\text{s}$ latency 4 mw asynchronous frame-free dynamic vision sensor using transimpedance preamplifiers," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 3, pp. 827–838, 2013.
- [59] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [60] T. J. L. Tan and R. Shokri, "Bypassing Backdoor Detection Algorithms in Deep Learning," in *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, Sep. 2020, pp. 175–183.
- [61] A. Tavanaei, M. Ghodrati, S. R. Kheradpisheh, T. Masquelier, and A. Maida, "Deep learning in spiking neural networks," *Neural networks*, vol. 111, pp. 47–63, mar 2019.
- [62] B. Tran, J. Li, and A. Madry, "Spectral Signatures in Backdoor Attacks," Tech. Rep., Nov. 2018, arXiv:1811.00636 [cs, stat] type: article.
- [63] V. Venceslai, A. Marchisio, I. Alouani, M. Martina, and M. Shafique, "NeuroAttack: Undermining Spiking Neural Networks Security through Externally Triggered Bit-Flips," in *2020 International Joint Conference on Neural Networks (IJCNN)*, Jul. 2020, pp. 1–8, iSSN: 2161-4407.
- [64] A. Viale, A. Marchisio, M. Martina, G. Masera, and M. Shafique, "Carsnn: An efficient spiking neural network for event-based autonomous cars on the loihi neuromorphic research processor," in *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021, pp. 1–10.
- [65] B. Wang, Y. Yao, S. Shan, H. Li, B. Viswanath, H. Zheng, and B. Y. Zhao, "Neural cleanse: Identifying and mitigating backdoor attacks in neural networks," in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 707–723.
- [66] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? a new look at signal fidelity measures," *IEEE signal processing magazine*, vol. 26, no. 1, pp. 98–117, 2009.
- [67] Z. Wang, A. C. Bovik, and L. Lu, "Why is image quality assessment so difficult?" in *2002 IEEE International conference on acoustics, speech, and signal processing*, vol. 4. IEEE, 2002, pp. IV–3313.
- [68] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [69] M. M. A. Weerasinghe, J. I. Espinosa-Ramos, G. Y. Wang, and D. Parry, "Incorporating structural plasticity approaches in spiking neural networks for eeg modelling," *IEEE Access*, vol. 9, pp. 117 338–117 348, 2021.
- [70] S. G. Wysoski, L. Benuskova, and N. Kasabov, "Evolving spiking neural networks for audiovisual information processing," *Neural Networks*, vol. 23, no. 7, pp. 819–835, 2010. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0893608010000924>
- [71] Y. Yao, H. Li, H. Zheng, and B. Y. Zhao, "Latent backdoor attacks on deep neural networks," in *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, 2019, pp. 2041–2055.
- [72] J. Yu, A. Lukefahr, D. Palframan, G. Dasika, R. Das, and S. Mahlke, "Scalpel: Customizing dnn pruning to the underlying hardware parallelism," *ACM SIGARCH Computer Architecture News*, vol. 45, no. 2, pp. 548–560, 2017.
- [73] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A survey of autonomous driving: Common practices and emerging technologies," *IEEE access*, vol. 8, pp. 58 443–58 469, 2020.
- [74] A. Zhang, X. Li, Y. Gao, and Y. Niu, "Event-driven intrinsic plasticity for spiking convolutional neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 5, pp. 1986–1995, 2021.
- [75] J. Zhang, C. Dongdong, Q. Huang, J. Liao, W. Zhang, H. Feng, G. Hua, and N. Yu, "Poison ink: Robust and invisible backdoor attack," *IEEE Transactions on Image Processing*, vol. 31, pp. 5691–5705, 2022.
- [76] A. Z. Zhu, D. Thakur, T. Özaslan, B. Pfrommer, V. Kumar, and K. Daniilidis, "The multivehicle stereo event camera dataset: An event camera dataset for 3d perception," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2032–2039, 2018.

APPENDIX A ADDITIONAL EXPERIMENTS

A. Datasets

We use four datasets: N-MNIST [50], CIFAR10-DVS [39], DVS128-Gesture [4], and N-Caltech101 [50]. We use N-MNIST, CIFAR10-DVS, and N-Caltech101 because their non-neuromorphic versions are common benchmarking datasets in computer vision for security/privacy in ML. N-MNIST is a spiking version of MNIST [37], which contains 34×34 60 000 training, and 10 000 test samples. An asynchronous time-based image sensor (ATIS) [53] captured the dataset across the 10 MNIST digits on an LCD monitor. The CIFAR10-DVS dataset is the spiking version of the CIFAR10 [34] dataset, which contains 9 000 training, and 1 000 test 128×128 samples,

corresponding to 10 classes. The N-Caltech101 dataset is the spiking version of the original Caltech101 [18] dataset. The dataset contains $180 \times 240 \times 8709$ train, and 823 test samples. To be consistent with the rest of the dataset, we square the data by cropping them to 180. Lastly, the DVS128-Gesture dataset is a “fully neuromorphic” dataset created for SNNs tasks. The DVS128-Gesture dataset collects real-time motion captures from 29 subjects making 11 different hand gestures under three illumination conditions, creating 1 176 128×128 training samples and 288 test samples. For all datasets, the samples’ shape is $T \times P \times H \times W$, where T is the time steps (we set it to $T = 16$), P is the polarity, H is the height, and W is the width.

B. Network Architectures

We consider three network architectures for the victim classifiers used in related works [17]. The N-MNIST dataset’s network comprises a single convolutional layer and a fully connected layer. For the CIFAR10-DVS dataset, the network contains two convolutional layers followed by batch normalization and max pooling layers. Then, two fully connected layers with dropout are added, and lastly, a voting layer—for improving the classification robustness [17]—of size ten is incorporated. Finally, for the DVS128-Gesture and N-Caltech101 datasets, five convolutional layers with batch normalization and max pooling, two fully connected layers with dropout, and a voting layer compose the networks. For more details, see our code repository.

Based on previous work [13], the spiking AE for the dynamic attack has four convolutional layers with batch normalization, four deconvolutional layers with batch normalization, and \tanh as the activation function for the DVS128-Gesture, CIFAR10-DVS, and N-Caltech101 datasets. For N-MNIST, we use two convolutional and two deconvolutional layers with batch normalization and \tanh as the activation function, which is the common AE [24] structure.

C. Default Training Settings

For training, we set a default learning rate (LR) of 0.001, MSE as the loss function, Adam as the optimizer, and we split the neuromorphic datasets in $T = 16$ frames using the SpikingJelly framework [16]. For the N-MNIST dataset, we achieve a (clean) accuracy of 99% on a holdout test set in 10 epochs. For the CIFAR10-DVS case, we achieved 68% accuracy after 28 epochs. For N-Caltech101, we achieve 76% accuracy after 30 epochs¹⁰, and a 93% accuracy with 64 epochs for the DVS128-Gesture dataset, see Table II. The results are aligned with the state-of-the-art [56].

Table II: Baseline training results for different datasets.

Dataset	# Epochs	Accuracy (%)
N-MNIST	10	99.4 ± 0.06
CIFAR10-DVS	28	68.3 ± 0.28
DVS128-Gesture	64	92.5 ± 0.91
N-Caltech101	30	75.5 ± 0.04

¹⁰Results for N-Caltech101 can be found in Appendix Section A-K due to space limitation.

D. Results for Static and Moving Backdoors

In Figure 10 and Figure 11, we provide results for static backdoors and moving backdoors, respectively.

E. Clean Accuracy Degradation of Static Triggers

The clean accuracy degradation after the static attack in the bottom-right corner is shown in Figure 12a and for the middle trigger in Figure 12b. Lastly, the degradation in the top-left corner is shown in Figure 12c.

F. Clean Accuracy Degradation of Moving Triggers

The clean accuracy degradation after the moving attack in the bottom-right corner is shown in Figure 13a and for the middle trigger in Figure 13b. Lastly, the degradation in the top-left corner is shown in Figure 13c.

G. Clean Accuracy Degradation of Smart Triggers

The clean accuracy degradation of the smart attack in the most active area and the least active trigger is shown in Figure 14a. For the least active area and the least active trigger, see Figure 14b. For the most active triggers in the most and least active areas, see Figure 15a and Figure 15b, respectively.

H. Additional Experimentation on STRIP

Figure 16 shows the normalized entropy for moving and smart triggers on different datasets.

I. Additional Experimentation on Pruning

Figure 17 shows additional experimentation solely on pruning—without retraining for a few epochs on different trigger types and datasets.

J. Additional Experimentation on SSIM

Figure 18 shows SSIM for smart triggers.

K. Experimentation on N-Caltech101

This section discusses the results of different attacks on the N-Caltech101 dataset.

1) *Static Backdoor*: For the N-Caltech101 dataset, we observe a similar behavior to the before-mentioned datasets. However, it is important to notice that the number of classes is ten times larger, making both main and backdoor tasks more challenging. Static triggers, see Figure 10, show that a large trigger, e.g., 0.1, is needed to inject the backdoor. However, the backdoor can also be effective with smaller triggers in some settings where the trigger is placed in the middle or the top-left corner.

2) *Moving Backdoor*: Moving trigger for the N-Caltech101 dataset, see Figure 11, shows similar performance and properties as in previously evaluated datasets. Similarly, large triggers are required to achieve high backdoor performance. The importance of the polarity selection is also evident in this case, where $p = 0$, i.e., background polarity, only works when the trigger is placed in the middle. This is because the image samples are centered, i.e., the image’s action area is in the middle.

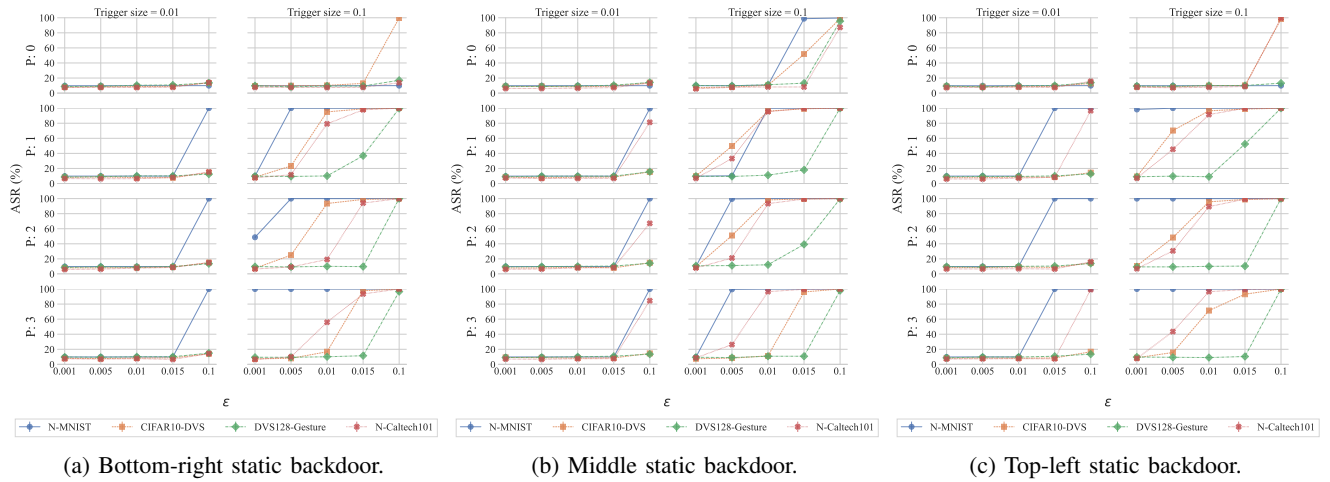


Figure 10: ASR of static triggers.

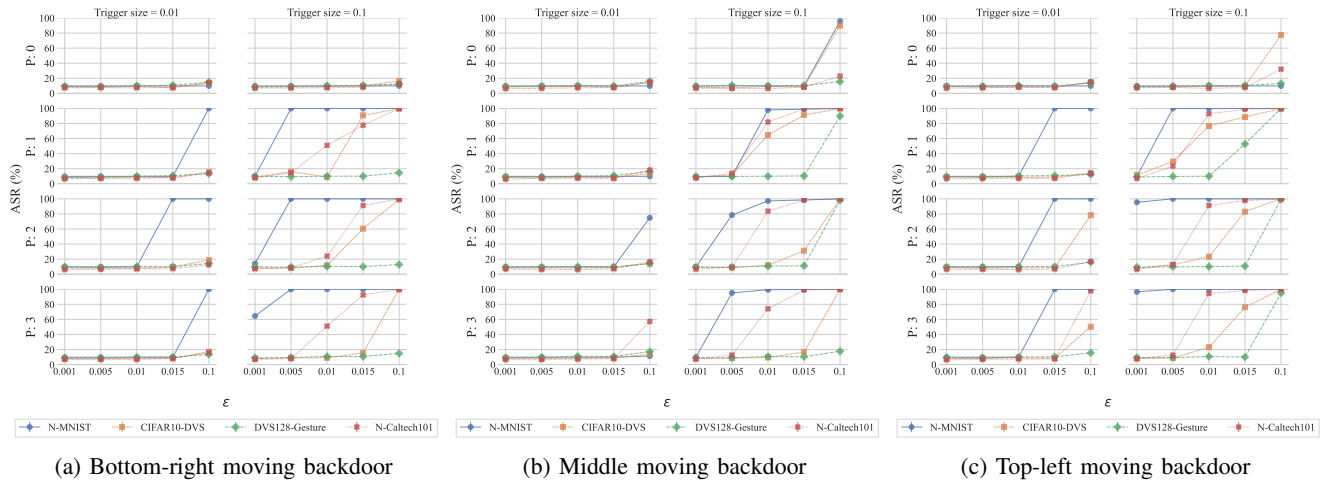


Figure 11: ASR of moving triggers.

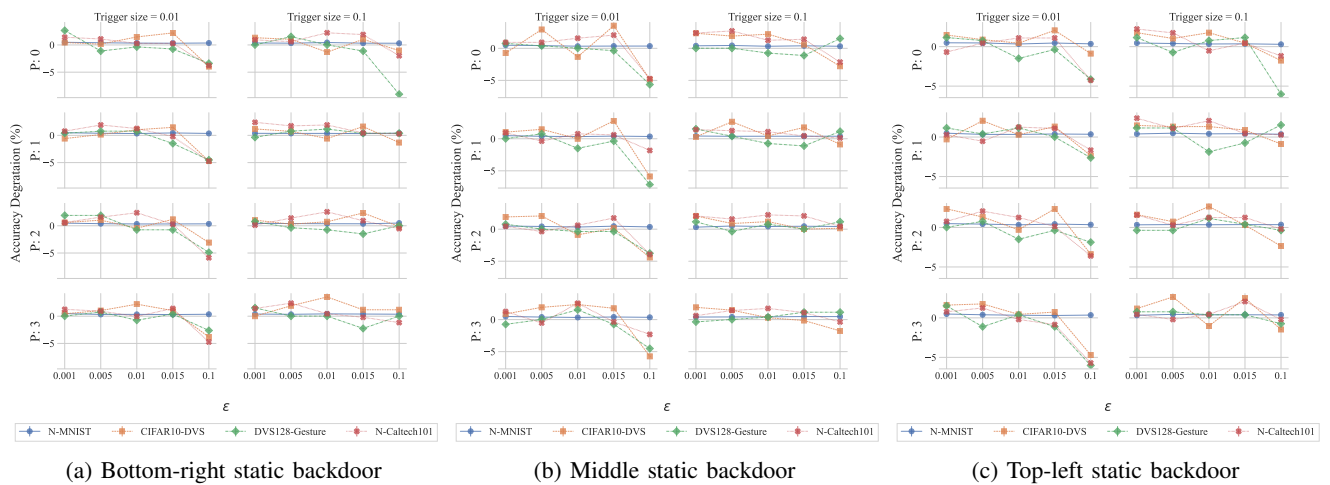


Figure 12: Clean accuracy degradation of static triggers.

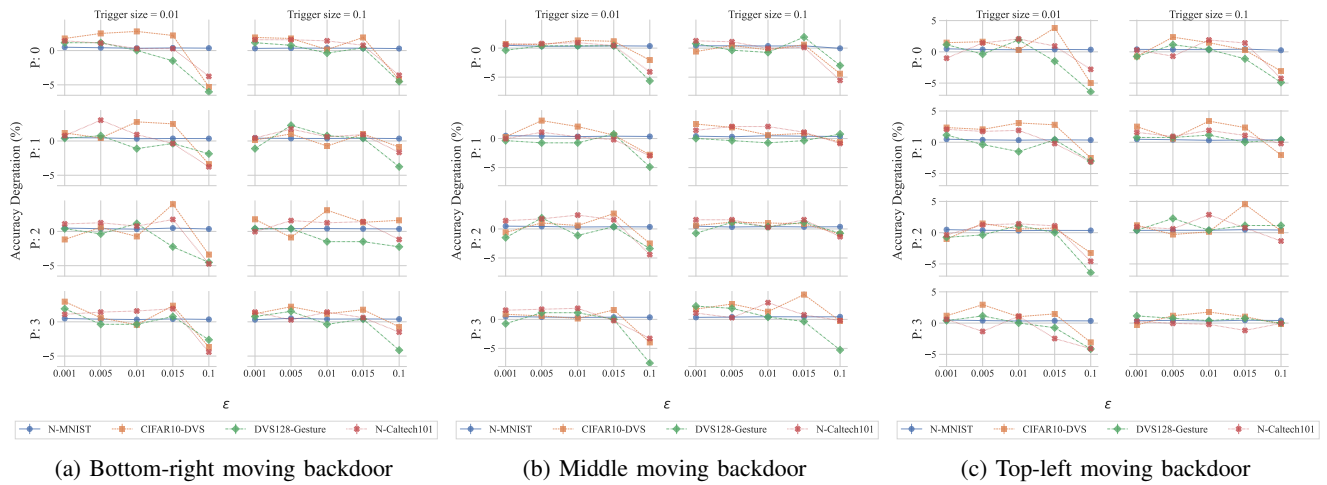


Figure 13: Clean accuracy degradation of moving triggers.

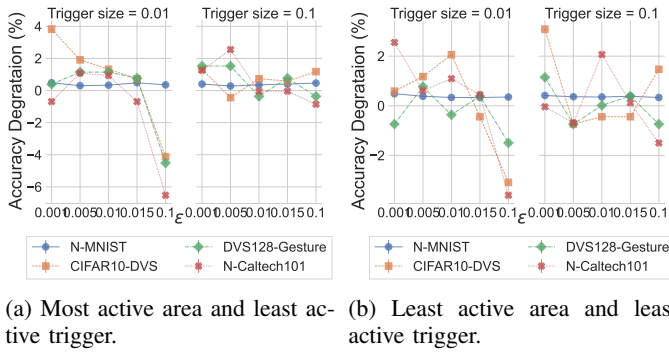


Figure 14: Clean accuracy degradation of the smart trigger in different settings.

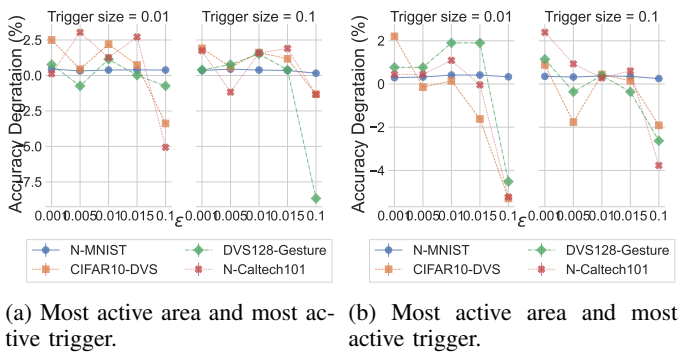


Figure 15: Clean accuracy degradation of the smart trigger in different settings.

3) *Smart Backdoor*: In smart backdoors, the trigger location and polarity are selected automatically. A large dataset like N-Caltech101 allows a better selection of these parameters because the location better represents the whole population, i.e., the mean is more accurate. Thus, despite being a “complex” dataset with many classes and noise, it achieves excellent results even when a simpler dataset cannot achieve a successful backdoor, see Figure 3d. Overall, as shown in Figure 1d, the results are aligned with the other datasets.

4) *Dynamic Backdoor*: Sample-specific dynamic triggers are also suitable for large datasets with many classes; see Figure 2. Note that having many classes makes the backdoor injection and the main task more challenging even though the attack can still achieve 100% ASR in most settings. Similar to the datasets analyzed before, when the trigger is less stealthy, i.e., $\gamma = 0.05, 0.1$, ASR is higher than in the most stealthy setup, i.e., $\gamma = 0.01$. Thus, there is a clear trade-off between stealthiness and backdoor performance. We observe a behavior similar to the use case with CIFAR-10, where $\alpha \geq 0.9$ exhibits the best performance.

APPENDIX B ARTIFACT APPENDIX

In this section, we include information about the reproducibility of the results and general usage of the artifact.

A. Description & Requirements

The repository¹¹ contains a README showcasing the different attacks and general information. Additionally, a *how to* guide is available. It contains detailed steps on how to prepare and download the datasets, as well as examples of how to run the code.

1) *How to access*: We provide access to the repository, including a guide on how to run the code. A permanent storage is also available¹².

¹¹<https://github.com/GorkaAbad/Sneaky-Spikes>

¹²<https://doi.org/10.5281/zenodo.10156889>

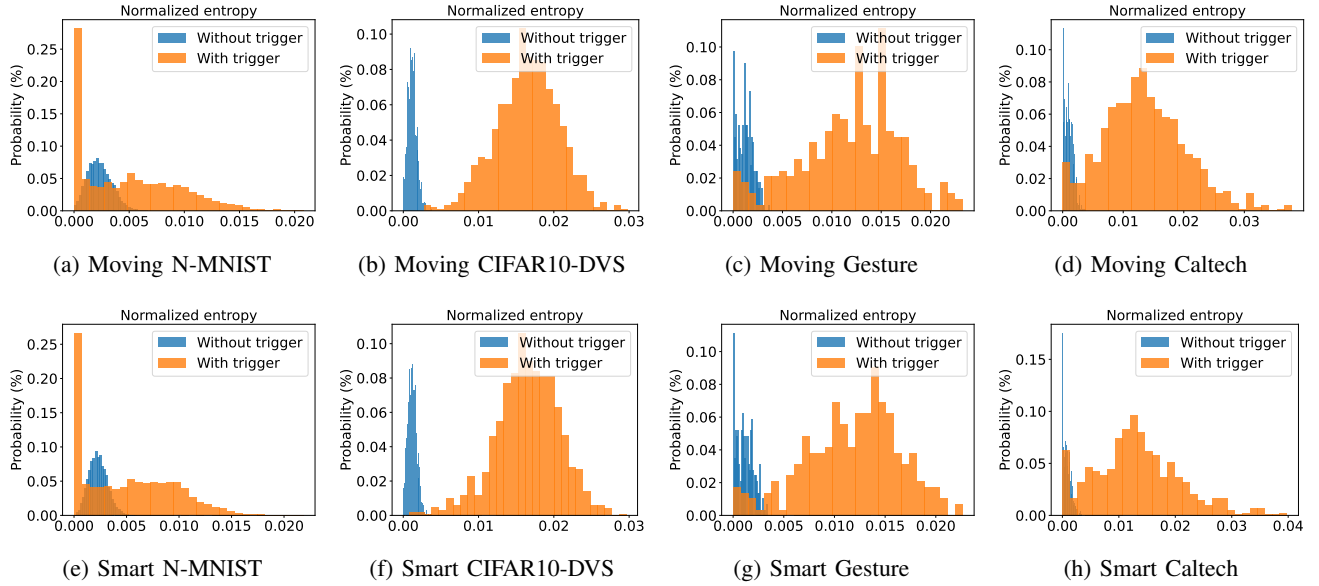


Figure 16: Normalized entropy of different triggers and datasets.

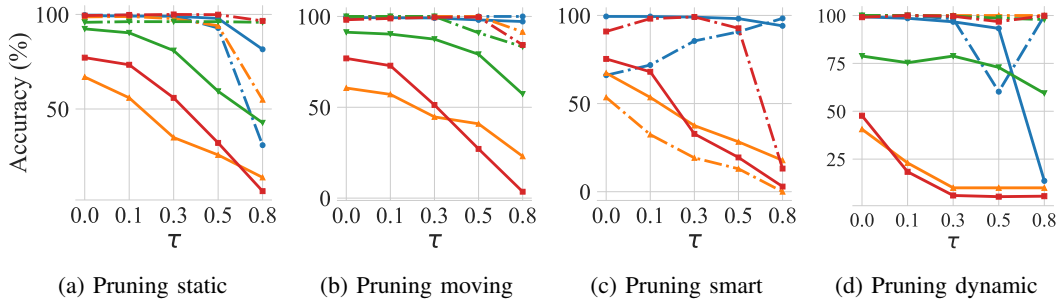


Figure 17: Effect of pruning on the ASR (dashed lines) and clean accuracy (full line) for different types of attacks, i.e., static, moving, smart, and dynamic. Blue corresponds to N-MNIST, orange to CIFAR10-DVS, green to the DVS128-Gesture, and red to N-Caltech101.

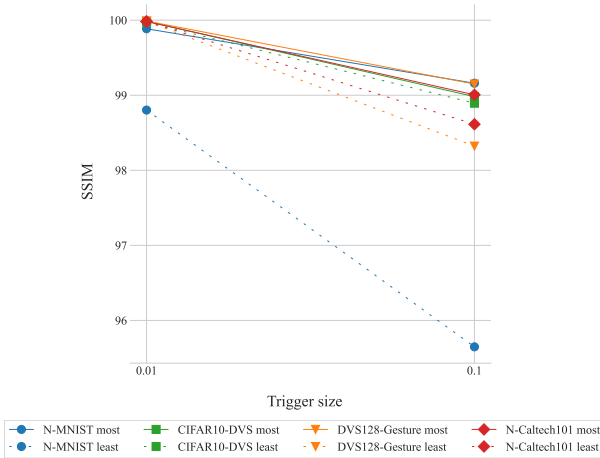


Figure 18: SSIM of smart triggers.

2) *Hardware dependencies:* To run the code, a GPU is strongly recommended. The code is tested on a machine with 1 NVIDIA A100 GPU with 40GB.

3) *Software dependencies:* The experiments were run on a Ubuntu 20.04 machine, using Python 3.8 and CUDA 11.7. We use the SpikingJelly¹³ framework to ease the training of SNNs and the processing of neuromorphic data.

4) *Benchmarks:* Some datasets are automatically downloaded. However, some others have to be downloaded manually. This is a restriction of SpikingJelly. To simplify the process of preparing the data, we provided a detailed *how to* guide in the repository.

B. Artifact Installation & Configuration

We require the standard libraries used in DL, e.g., torch, numpy, pandas. These can be easily installed

¹³<https://github.com/fangwei123456/spikingjelly>

with `pip`. We provide a guide on how to install the requirements in the repository.

C. Major Claims

This paper provides four types of attacks, i.e., static, moving, smart, and dynamic. Each is composed of a different type of backdoor trigger. We provide examples for reproducing the claims on the DVS128-Gesture and N-MNIST datasets (as they are the fastest to run). However, examples with additional datasets are available in the repository (in the `scripts/` folder).

- (C1): The static backdoor performance using a trigger of 10% the size of the input image, on the top-left corner, with a poisoning rate of 10% of the dataset, achieves around 100% ASR. This is proven by the experiment (E1) whose results are illustrated/reported in [Figure 10].
- (C2): The moving backdoor performance using a trigger of 10% the size of the input image, on the top-left corner, with a poisoning rate of 10% of the dataset, achieves around 100% ASR. This is proven by the experiment (E2) whose results are illustrated/reported in [Figure 11].
- (C3): The smart backdoor performance using a trigger of 10% using the most common polarity with a poisoning rate of 10% of the dataset, achieves around 100% ASR. This is proven by the experiment (E3) whose results are illustrated/reported in [Figure 3.B].
- (C4): The dynamic backdoor achieves around 100% ASR. This is proven by the experiment (E4) whose results are illustrated/reported in [Figure 4].

D. Evaluation

1) *Experiment (E1)*: [Static] [0 human-minutes + 1 compute-hour]: The static attack achieves around 100% ASR.

[How to] See [Execution] for executing the experiment.

[Preparation] Prepare the DVS128-Gesture dataset.

[Execution]

```
python main.py --dataset gesture --polarity 1
--pos top-left --trigger_size 0.1 --cupy
--type static --epsilon 0.1 --epochs 64
```

[Results] The results are stored in `results/`.

2) *Experiment (E2)*: [Moving] [0 human-minutes + 1 compute-hour]: The moving attack achieves around 100% ASR.

[How to] See [Execution] for executing the experiment.

[Preparation] Prepare the DVS128-Gesture dataset.

[Execution]

```
python main.py --dataset gesture
--polarity 1 --pos top-left
--trigger_size 0.1 --epsilon 0.1
--type moving --cupy --epochs 64
```

[Results] The results are stored in `results/`.

3) *Experiment (E3)*: [Smart] [0 human-minutes + 1 compute-hour]: The smart attack achieves around 100% ASR.

[How to] See [Execution] for executing the experiment.

[Preparation] Prepare the N-MNIST dataset.

[Execution]

```
python main.py --dataset mnist
--trigger_size 0.1 --epsilon 0.1
--type smart --most_polarity
--cupy --epochs 10
```

[Results] The results are stored in `results/`.

4) *Experiment (E4)*: [Dynamic] [0 human-minutes + 2 compute-hour]: The dynamic attack achieves around 100% ASR.

[How to] See [Execution] for executing the experiment.

[Preparation] Prepare the N-MNIST dataset.

[Execution]

```
python dynamic.py --dataset mnist
--cupy --epochs 10
--alpha 0.5 --beta 0.01
```

[Results] The results are stored in `results/`.

E. Notes

To execute the experiments for different types of attacks and datasets, we provided some scripts to ease the process. These can be found in the `scripts/` folder in the repository. Note that selecting the parameters during the execution is important and can affect the performance.