International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 13, Issue 4, 2025

Towards Resilient Neuromorphic Edge Computing: A Review of Mimicry Attacks and Federated Defence Strategies

[1]Shravan Kumar B, [2] Dr V. Naresh Kumar
BEST Innovation University

INTRODUCTION

Edge computing has changed real-time artificial intelligence (AI) applications by making it possible to make decisions and draw conclusions quickly at the edge of the network. At the same time, neuromorphic computing, which is based on the way that spiking neural networks (SNNs) function, has become a potential way to accomplish ultra-low-power edge inference in areas including autonomous systems, wearable health monitors, and smart IoT settings [1], [2].

Neuromorphic chips work in an event-driven and decentralized way, using dynamic synaptic weights, sparse connectivity, and asynchronous firing patterns [3]. This is different from typical von Neumann designs. This makes them very good at using energy in edge cases where there isn't much of it. But these same qualities also make them more vulnerable to a new type of hardware-level attack called Neuromorphic Mimicry Attacks (NMAs) [4]. NMAs take use of the non-deterministic, analog nature of neuromorphic systems to poison sensory inputs or change synaptic weights. They do this by pretending to be doing something normal while getting beyond standard intrusion detection systems (IDS) [5].

AI-based cybersecurity tools like anomaly detection and adversarial defense are becoming more popular in cloud and enterprise settings. However, they are very hard to use in edge networks because of the limited computing power and privacy risks that come with sending data [6], [7]. Federated Learning (FL) has become a way to protect privacy by allowing decentralized model training on edge devices without sending raw data [8]. FL frameworks, when used with homomorphic encryption, have shown good success in detecting IoT threats. For example, they were able to find more than 98% of DDoS assaults while using 20% less energy [9].

Even with these improvements, current FL-based cybersecurity frameworks don't deal with the particular threats posed by neuromorphic computing. One important thing to note is that traditional models are trained on inputs based on digital logic and can't capture the stochastic spiking behavior of SNNs or find small changes in neural networks. Also, existing FL systems presume that there are software-level vulnerabilities, but NMAs work at the physical and synaptic levels, changing the timing, weight distribution, and firing patterns of spikes in a way that isn't obvious [4], [5].

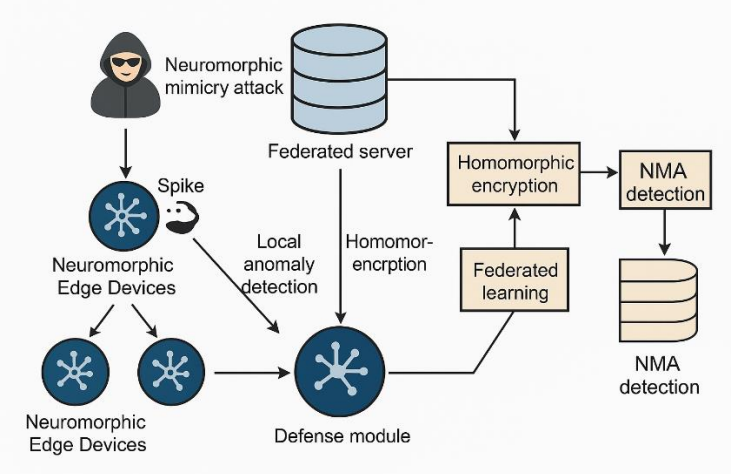

Fig. 1. A general idea of Federated-Adversarial Threat Defense in Neuromorphic Edge AI

This study suggests a new Federated-Adversarial Threat Defense Framework for Neuromorphic Edge AI to fill in these gaps. This system combines spike-sensitive anomaly detection with federated SNN training and secure model aggregation to find NMAs across a network of neuromorphic devices. This work makes important contributions, such as:

1. Creating a federated spike-aware anomaly detection model that takes into account changes in spike frequency, synaptic plasticity, and processing delay;
2. Using homomorphic encryption-based model aggregation to make sure that SNN learning is private and strong;
3. Neural-specific countermeasures, such as safe synaptic learning protocols and anomaly visualization for NMAs, have been put in place.
4. Validation using a proprietary neuromorphic dataset that included simulated imitation assaults showed meaningful gains in detection performance, communication efficiency, and explainability.

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 13, Issue 4, 2025

This study connects federated learning, neuromorphic computing, and adversarial defense. It makes it possible to deploy edge-AI in high-risk areas that are scalable, safe, and smart.

II. LITERATURE REVIEW

The combination of edge computing, neuromorphic architectures, and artificial intelligence (AI) has led to a lot of research into low-latency, real-time processing for smart cyber-physical systems. Neuromorphic threats and the problems with traditional intrusion detection systems (IDS), on the other hand, have shown that there are important research gaps in keeping edge AI settings safe.

Neuromorphic computing, which is based on how biological neurons work and how they are built, is becoming more popular in energy-efficient technologies like self-driving cars, wearable sensors, and decentralized control systems. These systems use spiking neural networks (SNNs) to handle input in a way that is both asynchronous and adaptable, which makes them perfect for edge AI situations [1], [2]. But their analog, probabilistic nature makes them more vulnerable in ways that are very different from how typical digital structures work.

Ravipati [1] came up with a new type of danger called Neuromorphic Mimicry Attacks (NMAs). These assaults use the low-level features of neuromorphic devices, including changing the weight of synapses and the frequency of spikes, to carry out hidden cyber intrusions that traditional IDS can't detect. NMAs work directly on the hardware learning dynamics, which makes neuromorphic edge AI systems less reliable than network-based or software-layer assaults.

Most current privacy-preserving cybersecurity solutions still depend on centralized or cloud-based systems, even though more people are becoming aware of these concerns. Rodríguez and Popescu [2] said that AI-powered security models based on homomorphic encryption, secure multiparty computation (SMPC), and federated learning are good at protecting privacy, but not much research has been done on how to use them in neuromorphic edge scenarios.

Also, current threat detection systems for edge AI generally put performance ahead of interpretability and don't think about how to run them quickly. Rahmati [3] stressed how important it is to build AI models that are easy to understand and work quickly for detecting cyber threats in real time at the edge. His suggested Explainable and Lightweight AI (ELAI) framework combines deep learning and federated learning with models that can be understood to find a compromise between accuracy and openness.

When used on neuromorphic systems, the problems with classical anomaly detection and IDS become clearer. These systems use sparse spike-timing-based communication, which doesn't work with digital signature-based detection. Ravipati [4], [5] showed that even little changes in the dynamics of synaptic weight or spike intervals can lead to big mistakes in system-level categorization without being noticed.

Zhong et al. [6] came up with EdgeShield, which uses attention-based lightweight detectors to protect edge AI against adversarial patch assaults. It works well in digital deep learning systems, but not in neuromorphic devices, which need spike-driven detecting techniques instead of static pixel-based protection.

Rupanetti and Kaabouch [7] wrote a thorough report on how AI-based countermeasures might help make IoT security better using edge computing. They stressed the need of combining decentralized trust models with collaborative threat detection, but they didn't apply these ideas to neuromorphic learning settings.

From a systems architectural point of view, Belcastro et al. [8] talked on how important it is to have safe service composition throughout the edge–cloud continuum. They suggested using multi-layered deployments that include near-edge, far-edge, and on-device computing. But their survey showed that there aren't any standardized or strong cybersecurity safeguards for neuromorphic processors that are part of these kinds of deployments.

Finally, Rahmati [9] suggested a Federated Learning-based cybersecurity architecture for IoT networks that uses homomorphic encryption and recurrent neural networks (RNNs) to find threats in real time while keeping data private. The framework performed a great job of finding DDoS and malware in regular IoT, but it didn't take into account the biological features or time-dependent behavior of neuromorphic SNNs. This led to more study.

In short, federated learning, edge AI, and anomaly detection have all made great progress, but there is still a huge gap in creating lightweight, neuromorphic-aware, federated threat prevention frameworks that can handle hidden low-level threats like NMAs. This study looks at this junction by developing a federated spike-aware architecture that can protect distributed neuromorphic edge nodes from being manipulated by enemies.

Table I: Comparative Analysis of Related Works in Neuromorphic Edge AI Security

Ref.	Authors / Title	Focus Area	Key Contribution	Limitations

[1]	H. Ravipati (2025)	Neuromorphic Security	Introduced *Neuromorphic Mimicry Attacks* (NMAs)	No defense model proposed; only theoretical attack modeling
[2]	Rodríguez and Popescu (2025)	Privacy-preserving AI in Cloud/Edge	Surveyed FL, HE, DP, and SMPC techniques for AI security	Lacks neuromorphic or spiking network-specific applications
[3]	M. Rahmati (2025)	Explainable Lightweight AI (ELAI) for Edge IDS	Proposed an explainable IDS with SHAP and attention-based deep learning	Not tailored for neuromorphic or spiking behavior
[4], [5]	H. Ravipati (2025)	NMA Impact & Dataset	Designed dataset simulating NMAs; analyzed spike frequency and latency effects	Does not address FL integration or distributed threat mitigation
[6]	Zhong et al. (2024)	Adversarial Patch Defense at Edge	Proposed *EdgeShield*, a lightweight patch-detection framework	Only supports image-space attacks; not applicable to SNN or NMAs
[7]	Rupanetti and Kaabouch (2024)	EC-IoT Security with AI	Reviewed AI-based countermeasures in edge-enabled IoT systems	No coverage of neuromorphic chips or federated anomaly models
[8]	Belcastro et al. (2025)	Edge–Cloud Continuum Survey	Defined architecture layers, deployment tools, and platform security trends	No mechanisms addressing neuromorphic vulnerabilities or adversarial defense
[9]	M. Rahmati (2025)	Federated Learning for IoT Security	Developed FL+RNN+HE architecture for privacy-preserving IDS	Focused on traditional IoT data (non-spiking); lacks neuromorphic compatibility

Research Gap

Even while federated learning, edge AI, and adversarial threat detection have all made progress recently, there are still important research gaps at the crossroads of neuromorphic edge computing and privacy-preserving cybersecurity frameworks:

1.There aren't any threat models for neuromorphic edge AI

Neuromorphic circuits provide low-power, real-time intelligence for edge applications, but their distinctive analog, asynchronous, and probabilistic behavior makes them less secure than regular digital systems. Neuromorphic Mimicry assaults (NMAs) show that typical digital IDS solutions are not enough, yet there is no complete threat prevention architecture for these kinds of assaults right now [1], [4], [5].

2.No Federated Defense for Spiking Neural Networks (SNNs)

Most federated learning (FL) frameworks made for IoT or cloud contexts are based on classic deep learning models that are trained on image or tabular data [2], [9]. But SNN-based neuromorphic systems create temporal, event-driven spike streams that need completely new encoding, training, and aggregation methods. None of these are currently used in FL-based anomaly detection architectures.

3.Secure Aggregation and Neuromorphic IDS Don't Work Together

Homomorphic encryption and differential privacy have been added to FL frameworks for cloud or regular IoT security [2], [9], however these methods don't work in real-time SNN-based edge contexts

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 13, Issue 4, 2025

where energy, bandwidth, and latency are very limited.

4.Current lightweight IDS do not work with SNN models

Models like EdgeShield and ELAI have showed promise for lightweight on-device adversarial detection [3], [6], however they rely on standard CNN or hybrid deep learning architectures. These can't be used on neuromorphic platforms like Intel Loihi or IBM TrueNorth since they need security methods that work with spikes and events.

No Collaborative Anomaly Detection for Distributed Neuromorphic Edge Systems In edge–cloud continuum deployments, many neuromorphic nodes can analyze data from autonomous cars, wearables, or smart grids on their own [7], [8]. But there is no current solution that uses federated SNNs or shared anomaly indicators to coordinate threat detection across distant neuromorphic endpoints.

Research Objectives

The main purpose of this study is to create and put into action a safe, lightweight, and federated threat detection framework that is made for neuromorphic edge AI systems that are open to Neuromorphic Mimicry Attacks (NMAs). The following specific goals help to reach this overall goal:

Objective 1:

Make a spike-aware adversarial threat model for neuromorphic edge systems and design it.

Objective 2:

Create a lightweight framework for federated anomaly detection that works on distributed neuromorphic edge devices.

Objective 3:

Make a prototype of a real-time, end-to-end federated adversarial defensive system that works in edge-cloud contexts and test it.

Proposed Methodology

We suggest a new Federated-Adversarial Threat Defense Framework that is made just for neuromorphic edge AI systems to fill in the research gaps we found. The technique is broken down into three main steps: (1) Threat Modeling and Spike-Based Feature Extraction, (2) Federated SNN-Based Anomaly Detection with Privacy Preservation, and (3) Global Aggregation and Defense Reinforcement. These parts are spread out over many distributed neuromorphic edge nodes and work together using a lightweight, secure federated learning architecture.

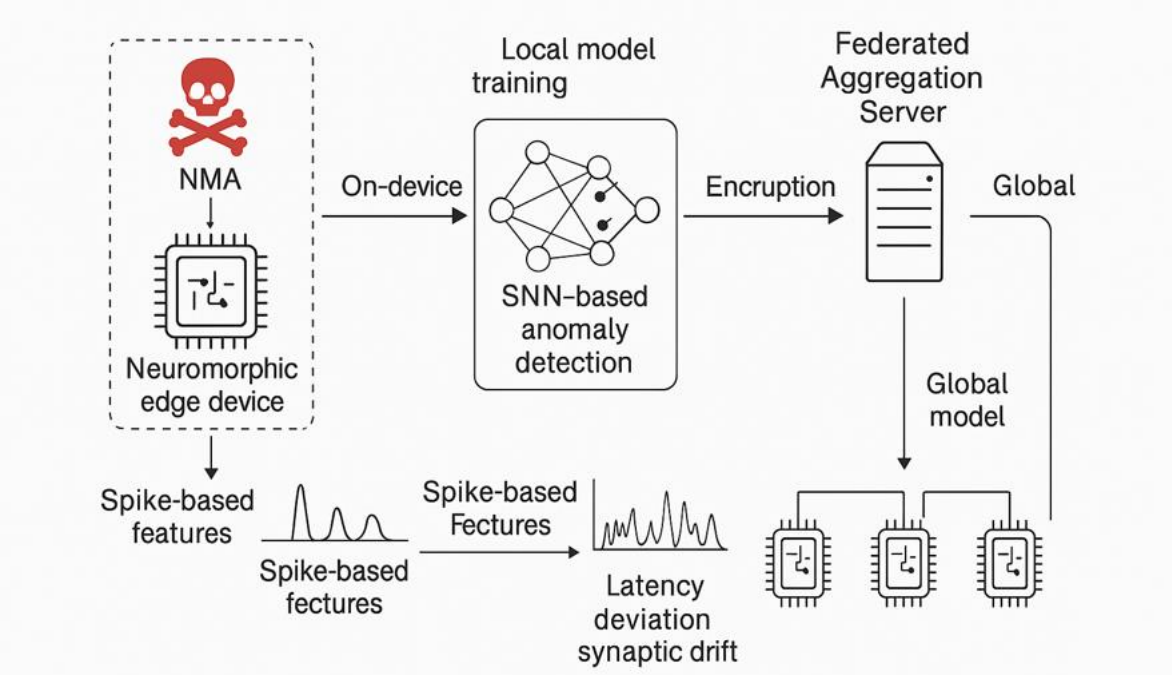

Fig. 2. Proposed approach that combines spike-based anomaly detection, federated SNN learning, and secure model aggregation to protect against NMAs.

1. Threat Modeling and Feature Extraction Based on Spikes

The initial step in the strategy is to describe Neuromorphic Mimicry Attacks (NMAs), which use the synaptic weights and spike frequencies in SNNs to change output decisions without leaving any clear evidence [1], [4]. We use simulated SNN environments like Brian2 or Nengo to add adversarial changes at the hardware or sensory input level and make fake attack samples. These samples help construct a neuromorphic threat dataset that includes important behavioral measures including spike frequency variation, latency deviation, and synaptic drift.

Each edge device has a neuromorphic processor (like the Intel Loihi) built into it that does spike encoding and feature extraction on the device itself. Metrics including average inter-spike interval (ISI), spike entropy, and frequency domain representations are

used to encode the local data into a time-series format. These parameters are used as input for localized anomaly detectors that can pick up on small changes in the brain's structure.

2. Using spike-compatible SNNs for federated anomaly detection

In the second step, each edge node gets a lightweight spiking neural network (SNN) that has been trained to find NMAs using the features that were taken out. These models work with sparse, event-driven data and are tuned for the neuromorphic processing architecture [5], [6], which is different from traditional IDS. Federated Learning (FL) is used by the system to safeguard data privacy and save network traffic. Each node trains its model privately without exchanging raw spike data.

After each local training epoch, each node uses homomorphic encryption or differential privacy noise injection [2], [9] to calculate encrypted model updates. Then, these encrypted updates (weights or gradients) are sent to a central aggregator for safe model fusion that protects privacy. This makes sure that even if edge nodes are hacked, sensitive patterns of activity on devices are not disclosed to outside actors.

3. GLOBAL AGGREGATION AND STRENGTHENING OF DEFENSE

In the third step, the Federated Aggregation Server collects encrypted model updates from all the neuromorphic devices that are taking part and uses them to create a global SNN model that is better able to handle different types of attacks. To make sure that the learning process is fair and honest, the server utilizes methods like secure model averaging and filtering out bad participants (for example, by using trust ratings or consensus voting).

The new global model is then sent back to edge devices, where it improves their ability to find anomalies in their own data. This feedback loop from the global level to the local level increases threat awareness across the system, even for zero-day assaults or localized abnormalities that may not have been seen before [3], [7]. Over time, this leads to a strong, cooperative neuromorphic security system that keeps changing to meet new threats in dynamic edge-cloud settings.

How it works:

1. Neuromorphic Mimicry Attack (NMA)

In fig 2, On the left side, a red skull icon represents a Neuromorphic Mimicry Attack (NMA). This attack targets neuromorphic chips—tiny processors in smart edge devices—by slightly altering their natural spike patterns, making them behave incorrectly without being detected by normal cybersecurity tools.

2. Spike-Based Feature Extraction

The neuromorphic edge device under attack continues to function, but it produces unusual spike patterns. These are shown in the form of wave-like plots labeled "Spike-based features" and include critical signs like:

- Abnormal firing rate (more or fewer spikes)
- Irregular time between spikes
- Hidden changes in how data flows through neurons (latency deviation and synaptic drift)

These subtle signs are used as clues to detect something is wrong.

3. Local SNN-Based Anomaly Detection

Each device has its own lightweight Spiking Neural Network (SNN) that runs locally to detect unusual behavior based on the extracted features. This step is labeled as “SNN-based anomaly detection”. It allows the device to spot problems on its own, without needing to send sensitive data to the cloud.

4. Secure Sharing through Encryption

Once the local model is trained or updated, it sends encrypted results to a central system. This is done using homomorphic encryption, which keeps all private data hidden—even while the server combines the results. The diagram shows this with an arrow labeled "Encryption" going to the Federated Aggregation Server.

5. Federated Aggregation Server

This server collects encrypted updates from many devices and creates a “Global Model”. This model is smarter because it has learned from many devices—each with its own version of what a normal or abnormal pattern looks like. The key idea is that devices share knowledge, not raw data, keeping everything private and secure.

6. Global Model Deployment

The improved global model is sent back to all neuromorphic devices. This makes every device stronger at spotting mimicry attacks, even if it hasn’t seen one before. In the diagram, this is shown with arrows from the global model to multiple neuromorphic chips at the bottom.

Conclusion:

In conclusion, this study presents a new federated defense architecture that is particularly made to find and stop Neuromorphic Mimicry Attacks (NMAs) in distributed edge AI settings. The suggested framework fixes major problems with current cybersecurity solutions that don't function with neuromorphic computing paradigms by combining spike-based anomaly detection, lightweight SNN models, and secure federated learning with homomorphic encryption. The technology guarantees real-time responsiveness, data privacy, and flexibility across different edge nodes. This is a big step forward in the safe use of spiking neural networks. This study sets the groundwork for neuromorphic security systems that can grow and adapt to new hardware-level threats while still protecting privacy.

REFERENCES

ISSN 2347–3657

Volume 13, Issue 4, 2025

[1] H. Ravipati, "Neuromorphic Mimicry Attacks: Using Brain-Inspired Computing for Hidden Cyber Attacks," arXiv preprint arXiv:2505.17094, 2025.

[2] A. Rodríguez and E. Popescu, "Privacy-Preserving AI Models for Cloud and Edge Computing Security," Synergy: Cross-Disciplinary Journal of Digital Investigation, vol. 3, no. 3, pp. 1–19, 2025.

[3] M. Rahmati, "Towards Explainable and Lightweight AI for Real-Time Cyber Threat Hunting in Edge Networks," arXiv preprint arXiv:2504.16118, 2025.

[4] H. Ravipati, "Simulated Neuromorphic Threat Dataset and Experimental Analysis of NMAs," arXiv preprint arXiv:2505.17094, 2025.

[5] H. Ravipati, "Neural-Specific Anomaly Detection and Learning Protocols for Spiking Systems," arXiv preprint arXiv:2505.17094, 2025.

[6] D. B. Li, X. Chen, and C. Zhong Liu, "EdgeShield: A Universal and Efficient Edge Computing Framework for Robust AI," arXiv preprint arXiv:2408.04181, 2024.

[7] D. Rupanetti and N. Kaabouch, "Combining Edge Computing-Assisted Internet of Things Security with Artificial Intelligence: Applications, Challenges, and Opportunities," Applied Sciences, vol. August 2024, 14, no. 16, p. 7104.

[8] L. F. Marozzo, A. Orsino, D. Talia, and P. Belcastro Trunfio, "Navigating the Edge-Cloud Continuum: A State-of-Practice Survey," arXiv preprint arXiv:2506.02003, 2025.

[9] M. Rahmati, "Federated Learning-Driven Cybersecurity Framework for IoT Networks with Privacy-Preserving and Real-Time Threat Detection Capabilities," arXiv preprint arXiv:2502.10599, 2025.