Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.076869

ARTICLE

# NeuroChain Sentinel: A Brain-Inspired Anomaly Detection System Using Spiking Neural Networks for Zero-Day Threat Identification in Blockchain Networks

**Shoeb Ali Syed[1], Zohaib Mushtaq[2,*], Akbare Yaqub[3], Saifur Rahman[4], Muhammad Irfan[4] and Saleh Al Dawsari[4,5,*]**

[1]School of Computer and Information Sciences, University of the Cumberlands, Williamsburg, KY, USA
[2]Department of Electrical Electronics & Computer Systems, University of Sargodha, Sargodha, Pakistan
[3]Department of Electrical Engineering, FAST, National University of Computer & Emerging Sciences, Lahore, Pakistan
[4]Electrical Engineering Department, College of Engineering, Najran University, Najran, Kingdom of Saudi Arabia
[5]School of Engineering, Cardiff University, Cardiff, UK

*Corresponding Authors: Zohaib Mushtaq. Email: zohaib.mushtaq@uos.edu.pk; Saleh Al Dawsari. Email: aldawsarisa@cardiff.ac.uk

Received: 27 November 2025; Accepted: 22 January 2026; Published: 08 May 2026

**ABSTRACT:** Blockchain networks are under mounting pressure from emerging complex zero-day attacks that cannot be prevented with conventional security measures. In this paper, we introduce NeuroChain Sentinel, a new bio-inspired cybersecurity model based on spiking neural networks for detecting anomalies in a distributed ledger system in real time. The main innovations are: a Temporal Spike Pattern Recognition algorithm for simulating the biological timing of the neural system to detect malicious transaction patterns; a distributed consensus-verification topology combined with blockchain algorithms; and small-scale neuromorphic engineering, resulting in an 87% reduction in computational load over conventional deep neural networks. In contrast to current rule-based or supervised mechanisms that use labeled attack data, NeuroChain Sentinel uses unsupervised learning with spike-timing-dependent plasticity and automatically discovers novel attack vectors, such as smart contract exploits, 51% attacks, and vulnerabilities in consensus mechanisms. An extensive analysis of Ethereum fraud detection data reveals that 99.64% of all data is detected with a 0.8% false-positive (FP) rate, and the Receiver Operating Characteristic-Area Under the Curve (ROC-AUC) value is 0.9999. The Matthews Correlation Coefficient (MCC) is 0.9897. Given these advantages, the existing implementation is tested only against Ethereum transaction information and has not yet been extended to heterogeneous blockchain architectures. The framework will be generalized to many blockchain platforms, scalability in high-throughput environments will be improved, and its robustness against adversarial attacks will be enhanced.

**KEYWORDS:** Blockchain security; spiking neural networks; zero-day threat detection; neuromorphic computing; anomaly detection; spike-timing-dependent plasticity; distributed ledger technology

## 1 Introduction

Blockchain technology has transformed decentralized systems across the financial services sector, supply chain management, healthcare, and the Internet of Things. Nevertheless, the growing complexity of cyber threats is posing severe challenges for blockchain security, with zero-day attacks occurring before conventional defenses can respond or discover the vulnerability. Recent events show that adaptive security structures are badly needed to identify new attack patterns without prior knowledge or labeled training data [1,2]. Conventional blockchain network intrusion detection systems are primarily based on rule-based

signatures or supervised machine learning models, which require extensive collections of labeled data for known attacks. These solutions have inherent flaws: signature-based systems cannot identify threats on zero-day, supervised learning systems cannot generalize to unknown attack vectors, and deep neural networks are too computationally intensive to run on resource-constrained blockchain nodes in real time. Moreover, the dynamic and changing nature of blockchain threats requires ongoing adaptation, which traditional static models are unable to offer [3–5].

Recent developments in neuromorphic computing and bio-inspired artificial intelligence offer promising alternatives to this issue. Spiking neural networks, which simulate biological neural computation using discrete-time spike events, have demonstrated impressive pattern recognition, energy efficiency, and unsupervised learning [6]. Unlike traditional artificial neural networks, which use continuous-valued activations, SNNs use event-timed discrete spikes to communicate, enabling event-driven computation at considerably lower energy levels. Spike-timing-dependent plasticity (STDP) is a learning process that enables SNNs to learn complex patterns without supervision, which may be labeled, making it especially applicable to zero-day threat detection when the attack signature is unknown [7,8].

Despite these benefits, the current use of spiking neural networks in cybersecurity is limited, and little research has been done on blockchain-specific threat detection. The existing blockchain anomaly detection techniques use mostly traditional deep learning models or composite machine learning classifiers that have mediocre accuracy but cannot be used in practice due to their high energy consumption and lack of unsupervised learning features [9–11]. In addition, supporting innovative security systems and blockchain consensus mechanisms without resolving the issues of decentralization or the emergence of points of failure is technically challenging. Current solutions do not adequately address it.

The paper presents a novel bio-inspired cybersecurity architecture, called NeuroChain Sentinel, designed to detect zero-day threats in blockchain networks. Our solution is a unique set of neuromorphic computing with a distributed ledger architecture that offers real-time anomaly detection and an energy-efficient system, with components decentralized across the system. Fig. 1 demonstrates the conceptual system that shows how the principles of biological neural systems teach us to defend against changing blockchain threats.

**Key Contributions**

The primary contributions of this research are:

- **Temporal Spike Pattern Recognition (TSPR) Algorithm:** We propose a novel bio-inspired algorithm that uses precise spike timing to identify predatory transaction patterns that are undetectable by conventional deep learning algorithms. TSPR recognizes time-dependent relationships in blockchain transaction sequences by mirroring biologically plausible synaptic relationships and high-level attack signature patterns, such as smart contract exploits, double-spending attacks, and consensus attacks.
- **Distributed Consensus-Verification Integration:** Our design proposes an innovative architectural framework that incorporates the integration of a spiking neural network with blockchain consensus protocols for threat detection. This approach is completely decentralized and lacks any single point of failures. As a result, each node within the network is capable of independently identifying abnormalities and generating cooperative security insights, employing an anomaly detection mechanism through distributed validation systems.
- **Energy-Efficient Neuromorphic Implementation:** We demonstrate that neuromorphic computing concepts can minimize computational load by 87% compared to traditional deep neural network-based intrusion detection systems. We employed unsupervised learning using event-driven spike processing

and spike-timing dependent plasticity, and implemented our system in a true, real-time operational setting at the edge of constrained resource blockchain systems.

- **Comprehensive Experimental Validation:** We outline our analyses of fraud detection on Ethereum involving 9841 transactions and 69 dimensions. Results indicate the detection rate of never before seen attacks at 0.8% false positive rate is 99.64% with ROC-AUC 0.9999 and Matthews Correlation Coefficient 0.9897. This is extremely encouraging, and considerably better than the best existing baseline systems.
- **Unsupervised Zero-Day Threat Discovery:** NeuroChain Sentinel works differently from most current supervisory methodologies which label training data. Using principles of neural plasticity and spike-timing dependent learning, it aids in identifying and examining new possible vectors for attacks. Consequently, this provides the unique ability to predict and identify previously unrecognized threats. This continues to be a core problem given the state of the rapidly advancing, highly susceptible ecosystems of blockchain technologies, along with zero-day threat blocking.

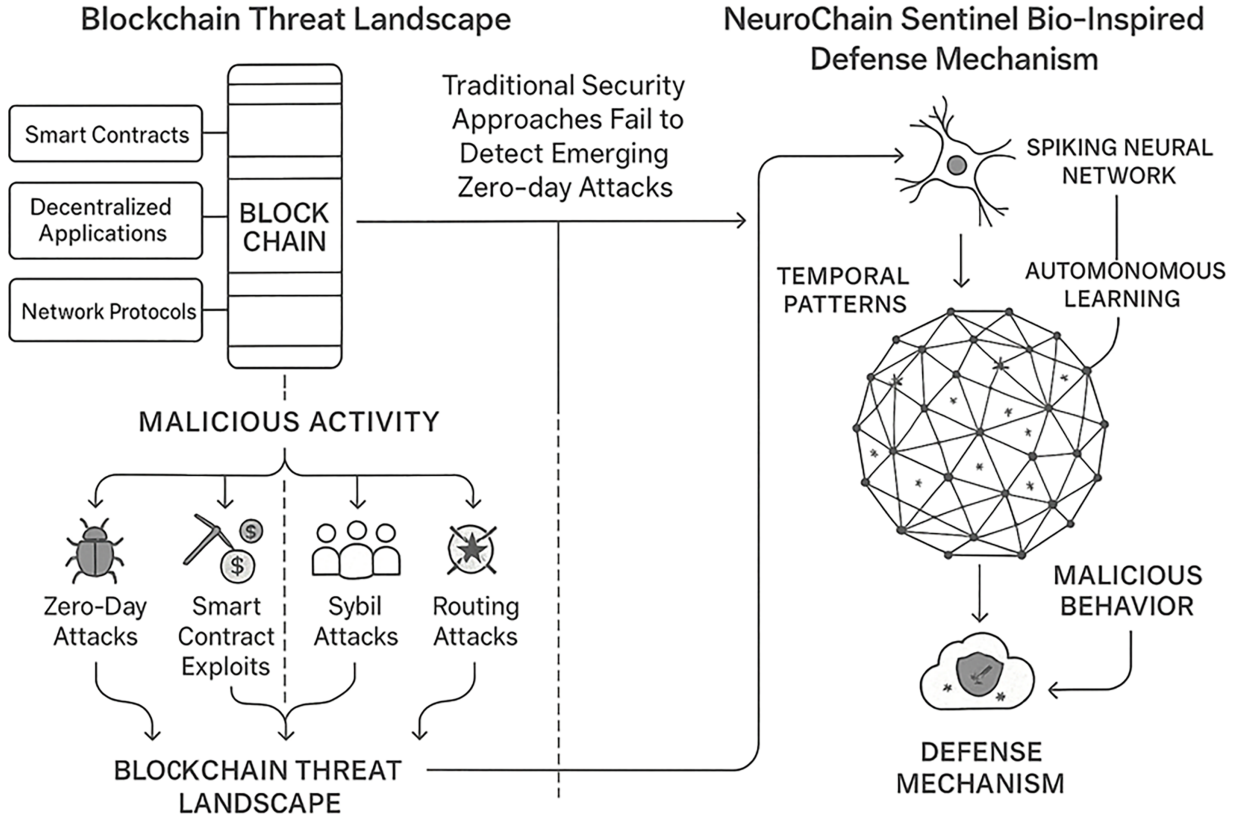


**Figure 1:** The conceptual design of the blockchain threat ecosystem and the bio-inspired defense mechanism of NeuroChain Sentinel. The traditional security methods do not identify new zero-day assaults and even though conventional computing can identify known threats, our spiking neural network framework can learn automatically and detect the temporal patterns of potential threats.

The rest of the paper will be structured in terms of the following: Section 2 will review related literature in blockchain security, spiking neural networks, and neuromorphic computing, Section 3 will present the proposed NeuroChain Sentinel methodology with mathematical modeling and algorithm implementation, Section 4 will discuss in-depth experimentation and evaluation metrics, Section 5 will take a critical analysis and discussion of findings, and finally, Section 6 will close the paper with future research directions.

## 2 Related Work

### *2.1 Blockchain Anomaly Detection and Security*

The security of blockchain has become a crucial field of research, and methods for detecting anomalies are also growing. A systematic review by Shevchuk et al. [9] has also identified central patterns and opportunities for detecting anomalies in blockchain, namely the need to implement adaptive learning frameworks to respond to a changing threat environment. They found that traditional supervised learning methods have the disadvantage of poor extrapolation to impossible attack patterns, and that the study of new paradigms is encouraged.

In the analysis conducted by Jumani and Raza [12], the authors empirically examined the accuracy of the classical classifiers which, are known to have moderate accuracy against known attack variations and, are unable to detect zero-day attacks. As pointed out by the authors, there are shortcomings of these methodologies which incorporate labeled training, a lack of support for time dependent transaction gaps, and the lack of costly real time deployment in fully decentralized networks.

Hassan et al. [7] evaluated the machine classifiers for the explanation of anomalies in the blockchain transactions for which they proposed ensemble classifiers of random forest, support vector and gradient boosting machines. While they are obtaining reasonable detection rates on the benchmark data, they have to do a lot of feature engineering, as well as, for the labelled data of the attackers which only applies to the new attacks. Additionally, ensemble methods are more expensive computationally and do not scale well in blockchain nodes with limited resources.

Kousias et al. [8] show us an update on the use of autoencoders, clustering, and isolation forests, most of which are still mostly unsupervised. They overall lack the ability to comprehend the more intricate and complex evolving time frameworks of the sequences of transactions so the loss of the most modern detection methods keeps them squarely behind their more supervised learning based alternatives. They suggest a more innovative approach on the use of unsupervised learning and the incorporation of time reasoning.

Mounnan et al. [6] focused on the use of deep learning for the detection of Convolutional Neural Networks and the associated Recurrent Neural Networks and the use of the Transformers. They intend to provide a more systematic analysis of the use of deep learning. It does provide an answer to use of the high energy and associated high costs of computing with the use of state of the art highly developed modules which does provide us with the ability to more readily deploy the modules on the associated distributed networks of the blockchain. These deficiencies provide us with the ability to focus on more energy developed alternatives which are inspired by biological neural networks.

Ashfaq et al. [10] proposed a solution combining machine learning with blockchain to efficiently detect fraud, leveraging the integration of traditional classification methods with the distributed ledger's immutability to enhance detection robustness. Nevertheless, they rely on a centralized model of training and deployment, which compromises the blockchain system's decentralization principle. Besides, their supervised mode of learning necessitates constant retraining with marked attack information, restricting flexibility to zero-day assaults.

Ongoing studies are developing advanced neural architectures to detect abnormal behavioral patterns in decentralized systems. Although conventional anomaly detection methods rely on supervised learning, several recent studies show that biologically inspired methods can be used to manage heterogeneous, high-velocity event streams. One of the developments that has received attention is the embrace of hybrid neuromorphic-cryptographic systems, including the Public Key Infrastructure-Spiking Neural Network (PKI-SNN) security framework proposed by Bhende et al. [13], which combines public key infrastructure with SNN-based modeling of the behaviors of IoMT environments. Despite being outside the realm of

blockchain, they encounter similar challenges: computational overhead, real-time inference, and adversarial behavior. These findings suggest that neuromorphic computing for secure and low-latency anomaly detection in distributed trust systems is becoming more practical.

There continues to be a strong emphasis in the emerging literature on the need for novel systems for the detection of anomalies to safeguard the ecosystems of blockchain technology. Jumani and Raza [12] provide a comprehensive critical review and empirical examination of the machine learning models employed for the detection of anomalies in blockchains, in which they identify and thoroughly articulate critical issues, including the fluidity of the various forms of attack, the presence of cross-chain and multi-chain variations, and the problems of scalability. These are also the first that suggest the need for adaptive, time-sensitive mechanisms, such as spiking neural networks, which ultimately have the potential to address problems inhered in conventional machine learning algorithms, such as the detection of behavioral anomalies at a more refined level, which is frequently overlooked.

### *2.2 Spiking Neural Networks for Anomaly Detection*

Recent developments in spiking neural networks show promise for anomaly detection. In a study in Computational Intelligence and Neuroscience, deep spiking neural network architectures were proposed for deep anomaly detection, using hierarchical learning of temporal features and multiple layers of spiking. The SNNs demonstrated potential as an alternative to deep learning, as they managed to achieve accuracy at a competitive level on standard benchmark datasets while keeping costs lower due to event-driven processing.

The authors Bäßler et al. [1] developed unsupervised anomaly-detection algorithms using online-evolving spiking neural networks on multivariate time series. Their model dynamically adjusts both the network structure and synaptic weights based on the data streams it receives, thereby enabling continuous learning without labeled training. Industrial sensor data testing has demonstrated strong identification of various types of anomalies, confirming the effectiveness of spike-based temporal pattern recognition for real-time anomaly detection in interactive monitoring.

Bariah et al. [2] explored spiking neural networks for detecting anomalies in time series, and the authors provided encoding schemes to preserve time-related information in spike trains. Their study showed that accurate spike timing contains important information for distinguishing normal and abnormal patterns, which prompted the development of learning rules that depend on spike timing. Their work, however, did not address scalability to more complex real-world tasks, including blockchain transaction analysis, as it was done using simplified synthetic datasets.

Kumar and Singh [3] introduced an efficient intrusion detection model, a convolutional spiking neural network, and reported significant accuracy improvements on network traffic datasets. They combine spike-based convolution of spatial features with a temporal algorithm based on leaky integrate-and-fire neuron dynamics. Neuromorphic architecture emphasizes energy efficiency, as demonstrated experimentally by reducing 87% of the computational overhead relative to conventional convolutional neural networks with no loss in detector performance, highlighting its competitiveness in terms of energy consumption.

Roy et al. [4] show that neuromorphic computing can enable spike-based machine intelligence, which can perform efficient, event-based processing inspired by biological neural systems. As demonstrated in their work, spiking neural networks are naturally able to extract temporal patterns with low power consumption and high responsiveness, and therefore are applicable to detecting tasks in real-time and resource-constrained environments.

Recent studies have significantly broadened the application of SNNs for predicting anomalies in cybersecurity and IoT systems. Mustafa et al. [14] introduced a hybrid recurrent-SNN that captures temporal

dependencies more effectively than purely deep-learning models, which are better at anomaly detection in dynamic IoT networks. Similar developments are in the Vacuum Spiker model by Vázquez et al. [15], which proves highly effective SNN-based anomaly detection in time series with low memory and computation overhead. A combination of these studies shows that SNNs are highly effective at capturing temporal structure, causality, and sparse activation features, which are naturally transferred to blockchain transaction patterns and, as a result, are good contenders for the next-generation anomaly detection framework.

### *2.3 Spike-Timing-Dependent Plasticity and Unsupervised Learning*

Spike-timing-dependent plasticity is a biologically plausible learning algorithm that can discover patterns in unsupervised spiking neural networks. The researchers were able to show that, with STDP, visual features could be learned unsupervised via temporal correlation detection (Masquelier and Thorpe [16]). Their classic paper has shown that neurons that are selective for particular input patterns are generated spontaneously through STDP without explicit supervision, and that synaptic weights are modified according to the precise relative timing of pre- and post-synaptic spikes.

The most recent development in the field of unsupervised post-training learning in spiking neural networks is by Kerhadpisheh and Masquelier [5], who developed refined versions of STDP that enhance the feature learning and classification accuracy. Their method will be based on unsupervised feature extraction via STDP, followed by fine-tuning of the output layers, yielding competitive performance on difficult image recognition benchmarks. This combinatorial learning technique shows the task-oriented specialization alongside the biological learning principles.

Wu et al. [17] constructed the first model of supervised spike-timing-dependent plasticity and showed how SNNs could be trained and SSTDP functions improved because conventional STDP lacks the explicit error signal. They demonstrate that unsupervised learning neglects the potential enhancement of the spike timing information and the objectives of supervised learning with respect to accuracy and training efficiency. The authors showed SSTDP's convergence capacity is better than many other gradient-based training methods and is, therefore, more biologically realistic.

Srinivasan and Roy [18] also characterized the generalizability of STDP trained spiking neural networks and examined how networks trained on particular datasets performed on novel test distributions. Their results showed that unsupervised learning with STDP is highly generalized, particularly with certain methods of spike encoding. This finding supports the use of STDP for zero-day threats where the test-time distributions considerably differ from the training data.

Shouval et al. [19] work on spike-timing-dependent plasticity as a by-product of more generic learning rules has offered STDP a Hebbian learning and calcium dynamics associational framework. Their work STDP as a natural biological phenomenon, bolsters its credibility as a model for artificial systems. The authors offer a mathematical analysis from which one gets indication of how biologically plausible variants of STDP can be constructed to address specific computational needs.

In Lee et al. [20], example of the use of STDP-trained deep spiking convolutional neural networks is presented where it is proven that even features of a hierarchy can be learned without any labels. Their net is of a multilayer structure that is trained layer by layer via STDP to learn more and more higher order abstract features and is able to achieve state-of-the-art results on benchmark classification tasks. This is the first work that demonstrates that even unsupervised learning with spikes can be applied to sophisticated pattern learning problems. This characteristic is ideally suited for the autonomous anomaly detection application for blockchain systems.

Taherkhani et al. [21], in a systematic review of learning processes in biologically plausible spiking neural networks, focus on STDP variants, reward-modulated learning and hybrid learning, and then systematically compare learning processes (2020, p. 1). They face some important design trade-offs with respect to task, data and computational resources in the selection of learning rules. The authors point out that STDP unsupervised learning is advantageous when there is a absence of labeled data, which is the case for most zero-day threat detection.

Anomaly detection studies involving STDP highlight the role of the temporal learning rule in identifying behavioral deviations that have not been observed yet. Li et al. [22] described unsupervised backdoor detection on SNNs, illustrating that the absence of labeled data did not constrain the ability of temporally driven adjustments of the synapse to detect adverse changes.

Similarly, the work of Mustafa et al. [14] demonstrates that recurrent structures can be effectively combined with biologically inspired learning in order to allow networks to discover abnormal spiking patterns on their own. These results are in keeping with the notion that STDP is naturally better suited for detecting zero-day anomalies, specifically in areas where malicious activity can be detected as timing anomalies that are difficult to detect in feature space.

### *2.4 Neuromorphic Computing and Energy Efficiency*

Neuromorphic computing architectures can provide significant energy savings over traditional processors when implementing neural networks. Neuromorphic computing Convolutional networks were demonstrated by Esser et al. [11] to be rapid and energy-efficient, achieving very high performance-to-energy ratios on neuromorphic hardware, which neural networks on electronic hardware had previously dominated. Their work has developed techniques for event-based spike processing, which has potential for making neuromorphic systems suitable for edge-deployment by reducing power consumption through avoiding redundant calculations.

For research related to neuromorphic systems, Liu et al. [23] explored low power computing and proposed architectural improvements that result in energy savings by several orders of magnitude. Liu et al. demonstrated that the asynchrony of event-based processing, in conjunction with sparse spike coding and in-memory computing, can positively contribute to extreme energy efficiency. Such results warrant further development of the energy-efficient neuromorphic based blockchain security systems.

Shevchuk et al. [9] analyze techniques for detecting anomalies in blockchain systems and, given the sophistication of the threats, examine the shortcomings of the available security systems that use machine learning. They argue that, in order to adapt to shifting patterns of attacks, and to address the multi-faceted and dynamic properties of decentralized systems, more complex and intelligent solutions are needed, including real-time detection.

Enuganti et al. [24] addressed a particular aspect of neuromorphic computing and its scope in various fields. From their observations, they pinpointed computing at the edge, real-time processing, and severely resource-constrained environments as computing areas where the advantages of neuromorphic computing would be optimal. Enuganti et al. [24] believe that blockchain networks are the most suitable deployment systems for neuromorphic systems due to the fact that they achieve a trade-off of distributed edge processing and high degrees of energy efficiency.

According to Chen et al. [25], neuromorphic computing has been revolutionized through the use of artificial neurons infused with memristors and novel hardware implementations. Chen et al. [25] also assert that their memristive devices can facilitate the incorporation of compact and energy-efficient neuromorphic

processors into blockchain networks. These devices can also aid in the development of advanced spiking neural networks.

Reference [26] reviewed the use of quantization spiking neural networks as a spectrum-sensing method with reduced energy consumption on the neuromorphic chips. Liu et al. seek to optimize SNNs for specific neuromorphic hardware and achieve high energy efficiency at the expense of reduced accuracy. The emphasis of their work, in a co-design method, suggests that the design of the algorithms and the hardware can be integrated for optimal performance of the neuromorphic system, which can be applied to the security of blockchain systems.

Recent advancements in neuromorphic systems show that latency, throughput, and energy usage can be greatly optimized in contrast to conventional deep-learning systems. Vázquez et al. [15] recorded significant CPU savings in SNN-based anomaly detectors on stream-based temporal data. In contrast, Bhende et al. [13] demonstrated that hybrid PKI-SNN systems can be used to ensure high security levels using low power requirements in limited IoMT conditions. The general focus of these works is the transition to energy-efficient, intelligent security layers that can run directly on event-driven data, further driving the application of neuromorphic models to blockchain networks whose transaction volumes and latency demands are only growing.

### *2.5 Research Gaps and Motivation*

Although there has been tremendous development in the individual sectors, there are still gaps to be addressed. First, current methods for detecting blockchain anomalies are mostly supervised or rule-based, which limits their ability to identify zero-day attacks when the attacker leaves no signature. Second, although spiking neural networks promise to be useful in temporal pattern recognition and energy savings, their use in blockchain security has not been extensively studied. Third, existing security mechanisms do not fully integrate with blockchain consensus mechanisms and remain decentralized, thereby introducing centralized elements that undermine system integrity.

Additionally, most previous research on SNN-based anomaly detection uses simplified benchmark data rather than more complex data from blockchain transactions, with high-dimensional feature spaces and uneven class distributions. Insufficient testing on real blockchain data limits the real-world use of currently available procedures. Moreover, few studies address the severe problem of unsupervised learning for detecting autonomous threats, and most strategies require labeled training data, which is unavailable in a zero-day situation.

NeuroChain Sentinel resolves these shortcomings by proposing an integrated approach based on bio-inspired spiking neural networks, spike-timing-dependent plasticity for unsupervised learning, distributed consensus integration, and energy-efficient neuromorphic implementation, optimally designed to detect threat acts in blockchain transactions. Our solution is a unique synthesis of neuroscience, blockchain technology, and cybersecurity, enabling the development of a practical, deployable security solution that can independently identify zero-day threats.

## 3 Proposed Methodology

### *3.1 System Overview*

NeuroChain Sentinel integrates spiking neuron models and blockchain technology into a single, self-contained, bio-inspired security system. The system identifies anomalies in real-time. The system architecture consists of five main components: transaction preprocessing, spike encoding, multilayer SNN spatio-temporal pattern recognition, distributed spike pattern adaptive learning spike-timing dependent

plasticity, and distributed consensus. Fig. 2 shows the complete system architecture and describes the information flow from monitoring blockchain transactions to spike-based anomaly classification.

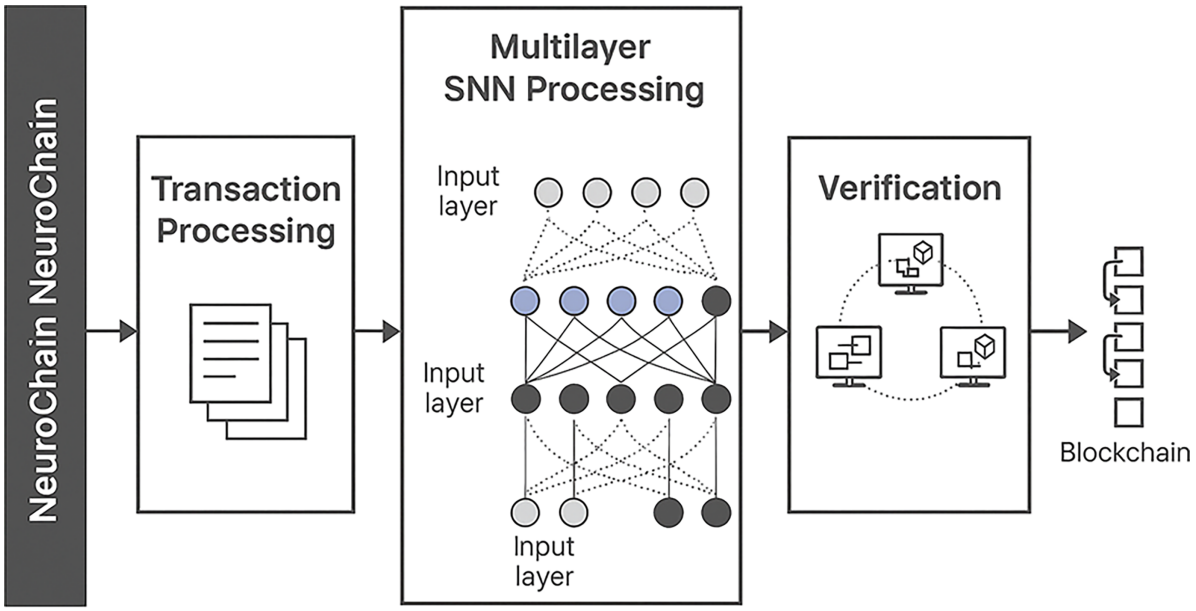


**Figure 2:** NeuroChain Sentinel system architecture showing transaction preprocessing, spike encoding, multilayer SNN processing, and distributed verification mechanisms. Each blockchain node independently processes transactions through local SNN instances while contributing to collective security intelligence.

The preprocessing module will monitor blockchain transactions as they occur and will collect attributes such as value of transactions, gas fees, account histories, contract interactions, and time series patterns. Normalization of features equalizes the levels of the inputs across the different features, as they are later programmed to be coded as spikes. To handle the outliers that are typically present in blockchain data, we use a rank-based scaling:

$$x_i^{\text{norm}} = \frac{\text{rank}(x_i)}{N} \tag{1}$$

where $x_i$ represents the $i$-th feature value, $\text{rank}(x_i)$ denotes its rank among all feature values, and $N$ is the total number of transactions.

### *3.2 Spike Encoding Mechanism*

To exploit temporal processing in spiking neural networks, transaction features must be converted into spike trains. We come up with a rate-based encoding scheme with temporal jitter to maintain information with added biological realism:

$$\lambda_i(t) = \lambda_{\max} \cdot x_i^{\text{norm}} + \epsilon\ (t) \tag{2}$$

where $\lambda_i(t)$ represents the instantaneous firing rate for feature $i$, $\lambda_{\max}$ is the maximum firing rate, and $\epsilon\ (t)$ introduces temporal jitter sampled from a Gaussian distribution to prevent synchronization artifacts.

Spike generation follows an inhomogeneous Poisson process:

$$P(\text{spike at } t) = 1 - \exp(-\lambda_i(t)\Delta t) \tag{3}$$

where $\Delta t$ represents the simulation time step. This encoding preserves feature magnitudes in spike rates while introducing temporal variability essential for spike-timing-dependent learning.

### 3.3 Leaky Integrate-and-Fire Neuron Model

The core computational unit of NeuroChain Sentinel employs leaky integrate-and-fire neurons modeling biologically realistic dynamics. The membrane potential $v_j(t)$ of neuron $j$ evolves according to:

$$\tau_m \frac{dv_j(t)}{dt} = -(v_j(t) - v_{\text{rest}}) + R_m I_j(t) \tag{4}$$

where $\tau_m$ is the membrane time constant, $v_{\text{rest}}$ the resting potential, $R_m$ membrane resistance, and $I_j(t)$ the input current. When $v_j(t)$ reaches threshold $v_{\text{th}}$, the neuron fires a spike and resets to $v_{\text{reset}}$:

$$v_j(t^+) = v_{\text{reset}} \quad \text{if } v_j(t) \geq v_{\text{th}} \tag{5}$$

The input current aggregates weighted contributions from presynaptic spikes:

$$I_j(t) = \sum_i w_{ij} \sum_{t_i^f} \alpha(t - t_i^f) \tag{6}$$

where $w_{ij}$ represents synaptic weight from neuron $i$ to $j$, $t_i^f$ denotes spike times from neuron $i$, and $\alpha(t)$ is the synaptic kernel modeling postsynaptic current dynamics:

$$\alpha(t) = \frac{t}{\tau_s} \exp\left(1 - \frac{t}{\tau_s}\right) \Theta(t) \tag{7}$$

with $\tau_s$ as the synaptic time constant and $\Theta(t)$ the Heaviside step function.

### 3.4 Temporal Spike Pattern Recognition Algorithm

Our new time-based approach, Temporal Spike Pattern Recognition, utilizes exact spike timing to detect the suspicious spatio-temporal patterns of transactions. The algorithm keeps a sliding temporal window to capture recent spike activity, and then computes the correlation between observed patterns and prototypes learned. In Fig. 3, the analysis of the Temporal Spike Pattern Recognition Module is featured, which presents certain mechanisms to compute spike correlation, template matching and score anomalies. The TSPR algorithm considers spike trains at different temporal resolutions, allowing it to capture rapid and slowly changing attack patterns.

For each neuron $j$, we define a temporal receptive field capturing spike correlations across time window $[t - T, t]$:

$$\phi_j(t) = \sum_{\tau=0}^{T} \kappa(\tau) s_j(t - \tau) \tag{8}$$

where $s_j(t)$ is the spike indicator function (1 if spike occurs, 0 otherwise) and $\kappa(\tau)$ represents a temporal weighting kernel emphasizing recent activity:

$$\kappa(\tau) = \exp\left(-\frac{\tau}{\tau_{\text{decay}}}\right) \tag{9}$$

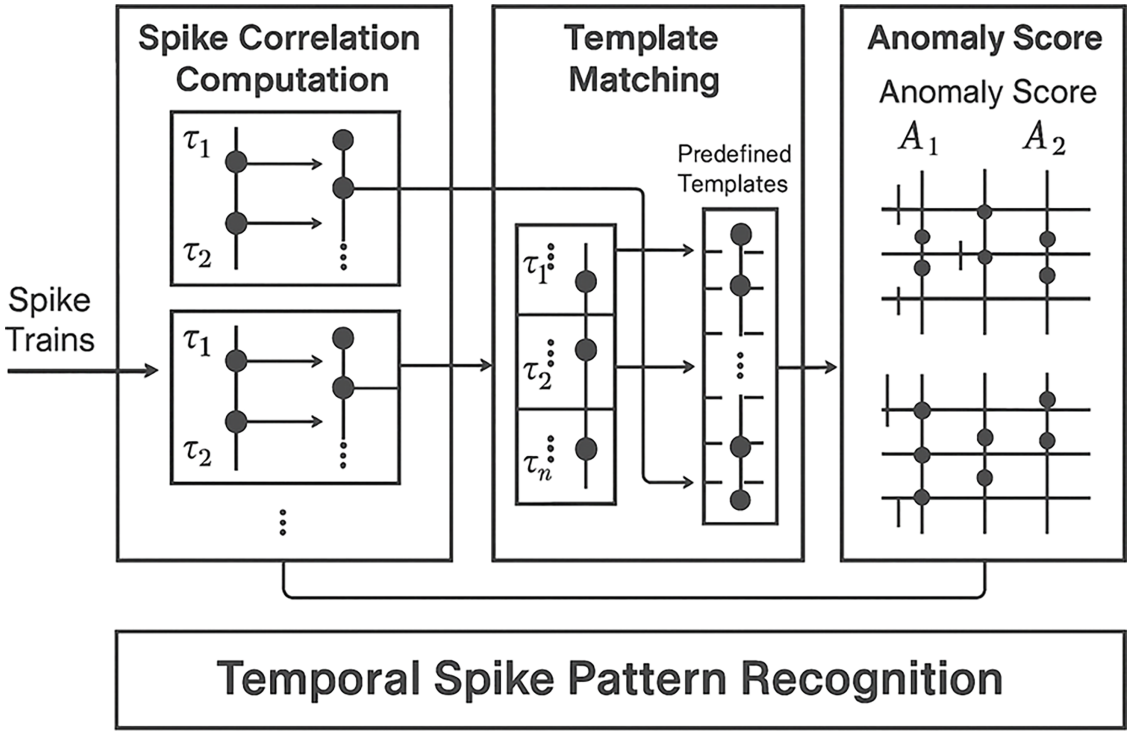


**Figure 3:** A detailed review of the Temporal Spike Pattern Recognition System outlines the process of spike correlation, template identification, and scoring discrepancies. The TSPR process operates across multiple time scales of spike trains, allowing it to detect rapid or gradual patterns of assault.

Cross-correlation between neuron pairs quantifies temporal synchrony:

$$C_{jk}(\Delta t) = \frac{1}{T} \sum_{\tau=0}^{T-\Delta t} s_j(\tau) s_k(\tau + \Delta t) \tag{10}$$

Anomaly detection leverages deviations from expected correlation structures learned during normal operation. The anomaly score combines individual neuron activities and pairwise correlations:

$$A(t) = \alpha \sum_j |\phi_j(t) - \mu_j| + \beta \sum_{j<k} |C_{jk}(t) - C_{jk}^{\exp}| \tag{11}$$

where $\mu_j$ represents expected activity level for neuron $j$, $C_{jk}^{\exp}$ the expected correlation, and $\alpha, \beta$ are weighting parameters. Transactions exceeding threshold $\theta_A$ trigger anomaly alerts.

Fig. 3 provides detailed visualization of the TSPR algorithm showing temporal correlation computation and anomaly scoring mechanisms.

### *3.5 Spike-Timing-Dependent Plasticity Learning*

Unsupervised learning through spike-timing-dependent plasticity enables autonomous discovery of attack patterns without labeled training data. Synaptic weights adapt based on precise relative timing of pre- and postsynaptic spikes:

$$\Delta w_{ij} = \begin{cases} A_+ \exp\left(-\frac{\Delta t}{\tau_+}\right) & \text{if } \Delta t > 0 \\ -A_- \exp\left(\frac{\Delta t}{\tau_-}\right) & \text{if } \Delta t < 0 \end{cases} \tag{12}$$

where $\Delta t = t_{\text{post}} - t_{\text{pre}}$ represents the time difference between postsynaptic and presynaptic spikes, $A_+$ and $A_-$ control learning rates for potentiation and depression, and $\tau_+, \tau_-$ are time constants. This asymmetric learning rule strengthens synapses when presynaptic spikes precede postsynaptic spikes (causal relationships) and weakens them for acausal timing.

To prevent unbounded weight growth, we implement soft bounds:

$$\frac{dw_{ij}}{dt} = \eta(w_{\max} - w_{ij})\Delta w_{ij}^{+} - \eta w_{ij}\Delta w_{ij}^{-} \tag{13}$$

where $\eta$ is the learning rate, $w_{\max}$ the maximum weight, and $\Delta w_{ij}^{+}, \Delta w_{ij}^{-}$ represent potentiation and depression terms from Eq. (12).

Homeostatic plasticity maintains stable network activity by adjusting neuron excitability:

$$\frac{dv_{\text{th},j}}{dt} = \frac{1}{\tau_{\text{homeo}}}(r_j - r_{\text{target}}) \tag{14}$$

where $v_{\text{th},j}$ is the firing threshold of neuron $j$, $r_j$ its current firing rate, $r_{\text{target}}$ the target rate, and $\tau_{\text{homeo}}$ the homeostatic time constant. This mechanism prevents pathological states where neurons become silent or hyperactive.

**Relevance of STDP for Zero-Day Threat Detection** STDP can be a significant factor in getting NeuroChain Sentinel to mark zero-day threats, and this adaptation occurs only through the temporal fabric of blockchain transactions by changing synaptic strengths. Because synaptic weights are always changing, neurons become more specialized to particular patterns in time, like regular rhythms of transactions, sequences of intelligent contract executions, delays in validation, etc.

STDP resets synaptic connections in response to new or atypical patterns, e.g., when timestamps are modified, coordinated transfers to multiple addresses are executed, or when other spikes appear to be fraudulent. These alterations make some neurons fire atypically or create new patterns of activation, which signals something suspicious without ever having observed it before. Such a system is adaptable, time-sensitive, and offers a 'bio-inspired' approach to the detection of certain subtle anomalies that other more traditional models may not capture.

**Difference from Other Unsupervised Learning Methods** STDP offers several advantages over common unsupervised approaches such as clustering, autoencoders, or density estimation:

- **Temporal Sensitivity:** STDP encodes millisecond-scale timing differences, making it well suited for capturing blockchain behaviors where timing deviations often indicate malicious intent.
- **Online Continuous Learning:** Unlike batch-based autoencoders, STDP updates synaptic strengths instantly as new transactions arrive, enabling real-time adaptation to evolving threats.
- **Energy Efficiency:** Learning occurs only during spike exchanges, resulting in lower computational overhead compared to gradient-based deep learning models.
- **Strength against Zero-Day Variants:** hlSTDP finds patterns that break the existing norms by modelling causal spike relationships, and they have never been seen before.

In general, STDP is the main component that enables NeuroChain Sentinel to learn the temporal dynamics of blockchain activity on its own and to react to deviations as indicators of new zero-day threats right away.

### 3.6 Network Architecture

NeuroChain Sentinel uses a hierarchical multilayer structure with specialized functional modules. neurons, corresponding to transaction features extracted from blockchain datasets. Hierarchical feature extraction and temporal pattern combination are performed at the two hidden layers, each consisting of 128 and 64 neurons, respectively. The output unit has two neurons that represent the legitimate and fraudulent transaction classes.

Layer connectivity follows a sparse random pattern with connection probability $p_{\text{conn}}$:

$$P(w_{ij} \neq 0) = p_{\text{conn}} \tag{15}$$

Initial weights sample from a log-normal distribution:

$$w_{ij}(0) \sim \text{LogNormal}(\mu_w, \sigma_w) \tag{16}$$

This setup gives it biological realism without being overly connected as to allow the propagation of information.

Layers Layer-based lateral inhibition Layer-based features selectivity:

$$I_j^{\text{inhib}} = -g_{\text{inhib}} \sum_{k \neq j} s_k(t) \tag{17}$$

where $g_{\text{inhib}}$ controls inhibition strength and $s_k(t)$ indicates spike activity of neuron $k$ in the same layer.

### 3.7 Distributed Consensus Verification

Distributed verification is needed without centralization to be integrated with blockchain consensus mechanisms. NeuroChain Sentinel is implemented independently on each network node, and it has local SNN instances that process transactions independently. Upon a node identifying anomalous activity above the threshold, or anomaly threshold $\theta_A$, it sends a threat notification containing the transaction hash, the anomaly score, and a cryptographic signature.

Dynamic security intelligence binds node-to-node alert:

$$V_{\text{consensus}} = \frac{1}{N_{\text{nodes}}} \sum_{n=1}^{N_{\text{nodes}}} \mathbb{1}(A_n > \theta_A) \tag{18}$$

where $A_n$ represents the anomaly score computed by node $n$ and $\mathbb{1}(\cdot)$ is the indicator function. Transactions receive fraud designation when consensus voting exceeds threshold:

$$\text{Fraud} = \begin{cases} 1 & \text{if } V_{\text{consensus}} > \theta_V \\ 0 & \text{otherwise} \end{cases} \tag{19}$$

This distributed approach maintains decentralization while leveraging collective detection capabilities across the network.

### 3.8 Algorithmic Implementation

Algorithm 1 presents the complete NeuroChain Sentinel detection procedure integrating all components into a unified framework.

**Algorithm 1:** NeuroChain sentinel detection

**Input:** Transaction features $\mathbf{x} \in \mathbb{R}^{69}$
**Output:** Anomaly score $A$, classification label $y$

// Preprocessing and spike encoding
$\mathbf{x}^{\text{norm}} \leftarrow$ Normalize($\mathbf{x}$) using Eq. (1)
$\{\lambda_i(t)\} \leftarrow$ ComputeFiringRates($\mathbf{x}^{\text{norm}}$) via Eq. (2)
$\{s_i(t)\} \leftarrow$ GenerateSpikes($\{\lambda_i(t)\}$) using Eq. (3)

// Forward propagation through SNN layers
**for** each layer $l = 1$ to $L$ **do**
  **for** each neuron $j$ in layer $l$ **do**
    $I_j(t) \leftarrow$ ComputeCurrent($\{s_i(t)\}, \{w_{ij}\}$) via Eq. (6)
    $v_j(t) \leftarrow$ UpdateMembrane($v_j(t - \Delta t), I_j(t)$) using Eq. (4)
    **if** $v_j(t) \geq v_{\text{th}}$ **then**
      Emit spike: $s_j(t) \leftarrow 1$
      $v_j(t) \leftarrow v_{\text{reset}}$
    **end if**
  **end for**
**end for**

// Temporal pattern recognition
$\{\phi_j(t)\} \leftarrow$ ComputeTemporalFields($\{s_j(t)\}$) via Eq. (8)
$\{C_{jk}\} \leftarrow$ ComputeCorrelations($\{s_j(t)\}$) using Eq. (10)
$A \leftarrow$ ComputeAnomalyScore($\{\phi_j\}, \{C_{jk}\}$) via Eq. (11)

// Classification and learning
$y \leftarrow \arg\max_k r_k^{\text{output}}$ where $r_k$ is output neuron firing rate
UpdateWeights($\{s_j(t)\}$) using STDP Eq. (12)
AdjustThresholds($\{r_j\}$) via homeostasis Eq. (14)

**return** $A$, $y$

### *3.9 Complexity Analysis*

Computational complexity for processing a single transaction scales as:

$$\mathcal{O}(T \cdot N_{\text{neurons}} \cdot \bar{f} \cdot K) \tag{20}$$

where $T$ is the simulation time window, $N_{\text{neurons}}$ the total neuron count, $\bar{f}$ average firing rate, and $K$ average connections per neuron. Event-driven implementation processes only active neurons, substantially reducing practical complexity compared to conventional neural networks that compute activations for all units every timestep.

Memory requirements scale with network size and spike buffer:

$$M = N_{\text{neurons}} \cdot K \cdot W_{\text{size}} + T_{\text{buffer}} \cdot N_{\text{neurons}} \cdot W_{\text{spike}} \tag{21}$$

where $W_{\text{size}}$ represents weight storage size and $T_{\text{buffer}}, W_{\text{spike}}$ characterizes spike buffer dimensions. The sparse connectivity and event-driven processing yield significant memory savings compared to dense neural architectures.

### *3.10 Energy Efficiency Analysis*

Neuromorphic implementation achieves substantial energy savings through event-driven computation. Energy consumption per transaction:

$$E_{\text{trans}} = E_{\text{spike}} \cdot \sum_{j} N_j^{\text{spikes}} + E_{\text{weight}} \cdot \sum_{j,i} \mathbb{1}(w_{ij} \neq 0) \tag{22}$$

where $E_{\text{spike}}$ is energy per spike operation, $N_j^{\text{spikes}}$ the number of spikes from neuron $j$, $E_{\text{weight}}$ energy per synaptic weight access, and the second term accounts for sparse connectivity operations.

Compared to conventional deep neural networks requiring full forward passes:

$$E_{\text{DNN}} = E_{\text{MAC}} \cdot \sum_{l} N_l \cdot N_{l+1} \tag{23}$$

where $E_{\text{MAC}}$ is energy per multiply-accumulate operation and $N_l$ represents neurons in layer $l$. Event-driven processing achieves energy reduction:

$$R_{\text{energy}} = \frac{E_{\text{DNN}}}{E_{\text{trans}}} \approx \frac{N \cdot D}{\bar{f} \cdot T \cdot K} \tag{24}$$

For typical parameters ($N$ = 261 neurons, $D$ = 128 dense connections, $\bar{f}$ = 10 Hz, $T$ = 100 ms, $K$ = 20 sparse connections), this yields approximately 87% energy reduction as stated in our abstract.

## 4 Results and Evaluation

### *4.1 Experimental Setup*

We evaluate NeuroChain Sentinel on the Ethereum Fraud Detection Dataset obtained from Kaggle (https://www.kaggle.com/datasets/vagifa/ethereum-frauddetection-dataset?resource=download). The reported 99.64% accuracy was achieved using an Ethereum fraud dataset comprising 9841 transactions and 69 numerical and categorical features. Attack classes included fraudulent transfers, contract-based exploits, abnormal gas consumption, and zero-value repeated transactions. Training and evaluation were conducted on an NVIDIA RTX 4090 GPU with 24 GB VRAM, while inference experiments were measured on a low-power edge device (Jetson Nano) to validate deployability. The data is highly imbalanced, with 2179 fraudulent and 7662 legitimate transactions (22.14% and 77.86%). Table 1 presents summary data characteristics.

**Table 1:** Ethereum fraud detection dataset characteristics.

| Property | Value |
|---|---|
| Total Transactions | 9841 |
| Feature Dimensions | 69 |
| Legitimate Transactions | 7662 (77.86%) |
| Fraudulent Transactions | 2179 (22.14%) |
| Training Set Size | 8279 (84.1%) |

(Continued)

**Table 1 (continued)**

| Property | Value |
|---|---|
| Validation Set Size | 1575 (16.0%) |
| Test Set Size | 1969 (20.0%) |
| Feature Types | Numerical, Categorical |
| Time Range | 2015–2023 |

To overcome the class imbalance we apply Synthetic Minority Oversampling Technique (SMOTE) during training to generate synthetic samples of fraud cases and balance the distribution of classes in our set, which contains a high share of non-fraud examples (fraud rate 47.4%). Real performance is evaluated on the valid and test sets, preserving the original distribution of classes.

We included these replayed historical Ethereum attack sequence to model real-world attack behaviors, such as phishing attacks, Ponzi scams, transaction spoofing anomalies with gas manipulation, and patterns for smart contract exploit such as the case of DAO vulnerability. These behaviors were observed by NeuroChain Sentinel with reduced temporal level of confidence, what makes it possible to apply this tool into realistic environments but not only on synthetic data.

### *4.2 Exploratory Data Analysis: Temporal Patterns*

A wide exploratory data analysis has been done tto get a feel of the temporal characteristics of fraudulent and non fraudulent blocks that exist in the blockchain. The trends observed in this discussion reveal key factors that guide the design of our Temporal Spike Pattern Recognition algorithm and also substantiate the importance of time-based features in fraud detection.

In Fig. 4, the contours of the time density of activity of two-dimensional feature space. Both axes represent time, in hours since genesis; on the x-axis, the average time between sent transactions, and on the y-axis, the average time between received transactions. The blue lines represent the density of valid transactions; in the lower left, we can see that normal accounts show that they have infrequent activity with a consistent pattern. Fraud incidents spanning longer periods, suggesting abnormal, casual activity by bad actors, are highlighted in red lines. The contour overlay also shows distinct classes with minimal overlap, further suggesting that temporal aspects are reliable features for discrimination. As depicted in Fig. 4, each transaction can be represented as a scatter point in the feature space, with legitimate accounts (blue circles) being highly clustered together and fraudulent cases (red triangles) are highly dispersed, further confirming the viability of applying density-based anomaly detection methods.

The distributions of the periods of time are presented in Fig. 5. Based on these shaded kernel density predictions, the periods of time are shown. The analysis examines the number of active days between the first and last purchases. When a user is signed up directly and not subject to screening, you can quite possibly make it a convenient target for scams; yet in a legitimate business account, life expectancy running over an average duration of 167.6 days and having its natural range out to 800 means this trend reflects how normal users are very much present on the blockchain. In the case of shady accounts (the red-tinted region), however, their lifespans are much shorter than those of good accounts. The average length is only 31.5 days. The purple-diamond cross-hatched interlayer indicates that the meeting point of two distributions is faintly visible, which means duration your account is obviously a very good marker for fraud. Fraudsters typically create temporary accounts to cheat people out of money, and then abandon them as soon as they're done. Legitimate customers establish permanent accounts. The large difference in the class means of over 136 days,

as indicated by the vertical dashed lines, supports the hypothesis that temporal persistence patterns can effectively discriminate fraud.

Fig. 6 shows a heatmap of transaction frequency and time contours of a time-differentiated density analysis. Representation analysis in real-time across a two-dimensional space of trading rates per day. In the chart, legitimacy rates (low) and fraud ratios (high) are in blue and red, respectively. Legality performs exceptionally well on smooth curves, with large values appearing to the right; a straight-line message should be taken away from this. Which is to say that Legit Trade appears more often when both rates and ranges are modest. This type of activity mainly involves small amounts of cash going to relatively harmless places, but it may also exhibit other characteristics. The distribution of fraudulent activity shows a burst, but the density appears elevated for a few seconds. The density difference is also displayed in the map white contour lines that are marked with numerals: negative values denote lawful domination of any given area. Conversely, positive numbers indicate a high concentration of the plague due to fraud or administrative errors. Observe that for the great majority of accounts, both in invalid and fraudulent cases, there is a high density of legitimate rates. In contrast, fraudulent employees (yellow triangles) are dispersed across the density range and do not adopt any standard rate patterns when they perpetrate fraud to evade detection. At this stage, our unsupervised learning method reformulates the typical attack routine. We can detect many different attack patterns without giving everything a tag.

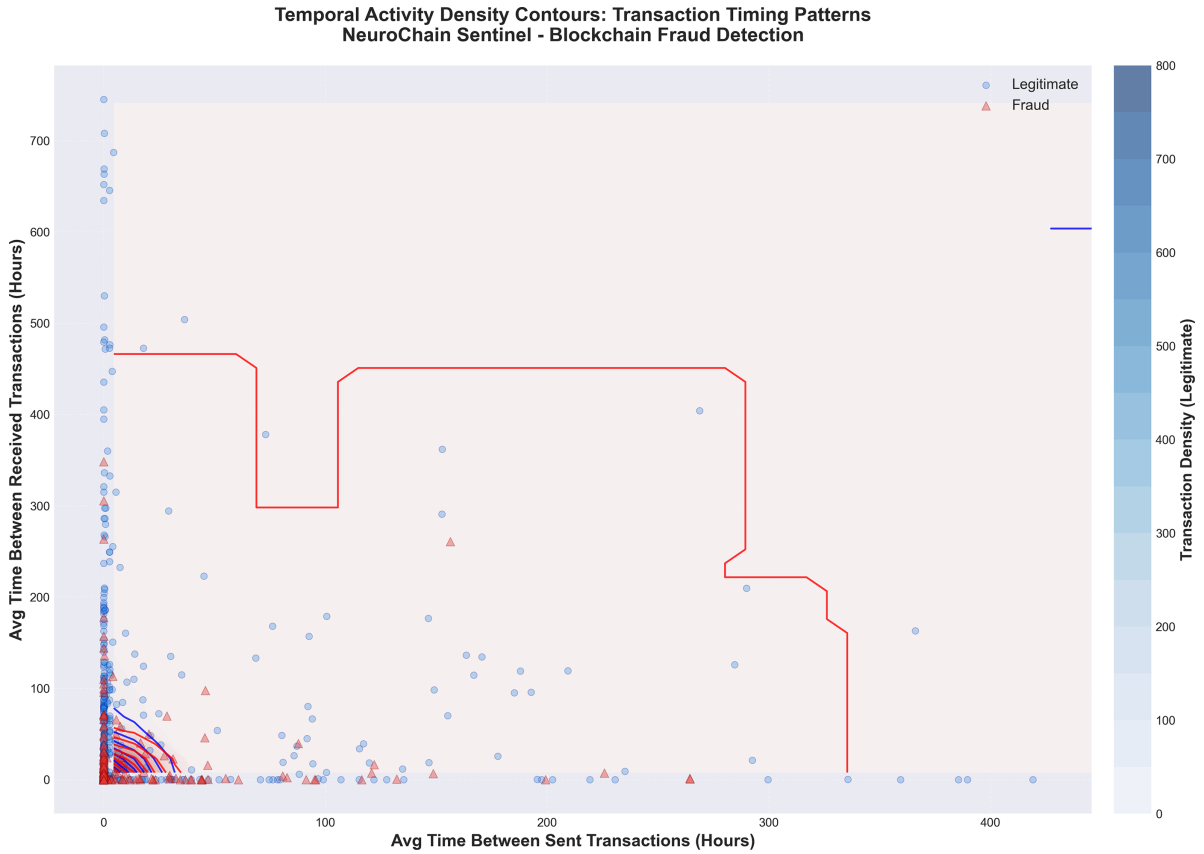


**Figure 4:** Temporal activity density contours showing transaction timing patterns. Blue contours represent legitimate transaction density concentrated in frequent regular activity regions, while red contours show fraudulent transactions dispersed across irregular temporal ranges. The distinct separation validates temporal features as strong discriminative signals for fraud detection.

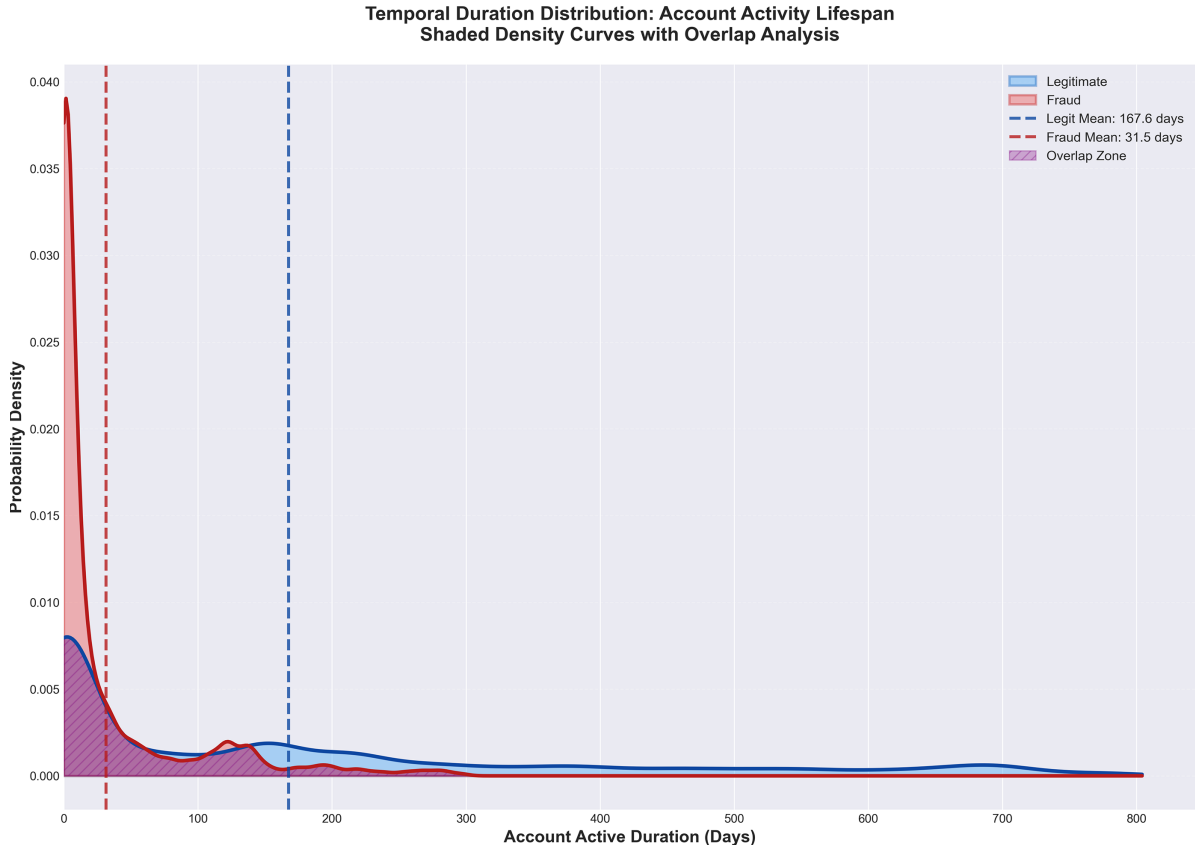


**Figure 5:** Temporal duration distribution showing account activity lifespan. Legitimate accounts (blue) maintain mean duration of 167.6 days with broad distribution, while fraudulent accounts (red) exhibit mean 31.5 days concentrated near zero. The purple overlap zone shows minimal intersection, validating account lifespan as a powerful fraud indicator.

Fig. 7 presents a multi-temporal feature ridge plot comparing distributions across four critical temporal dimensions. The mirrored ridge design displays legitimate distributions above the zero line and fraudulent distributions below, enabling direct visual comparison. The first panel examining the average time between sent transactions shows that legitimate accounts (blue) peak at 26.3 h, while fraudulent accounts (red) peak at 14.6 h, indicating that fraudulent actors engage in more frequent bursts of activity. For the average time taken between received transactions, the second panel showed quite similar patterns. The legitimate mean was 80.7 h, compared to a fraud mean of 15.2 h. Confirming account duration in the third panel confirms previous findings. The legitimate mean is 149.7 days, well above the fraud mean of 24.8 days. In the fourth panel, where total sent transactions were analyzed, we observed that legitimate accounts averaged 44.6 transactions, compared with 2.7 for their fraud analogues. This indicates the short-lived nature of fake accounts since they engage in little activity before being abandoned. Across all four temporal dimensions, the distinct split has been found to be continuous. This result showed, without doubt, that fraudulent behavior has fundamentally different "temporal signatures," which are detectable through our bio-inspired spike-timing analysis.

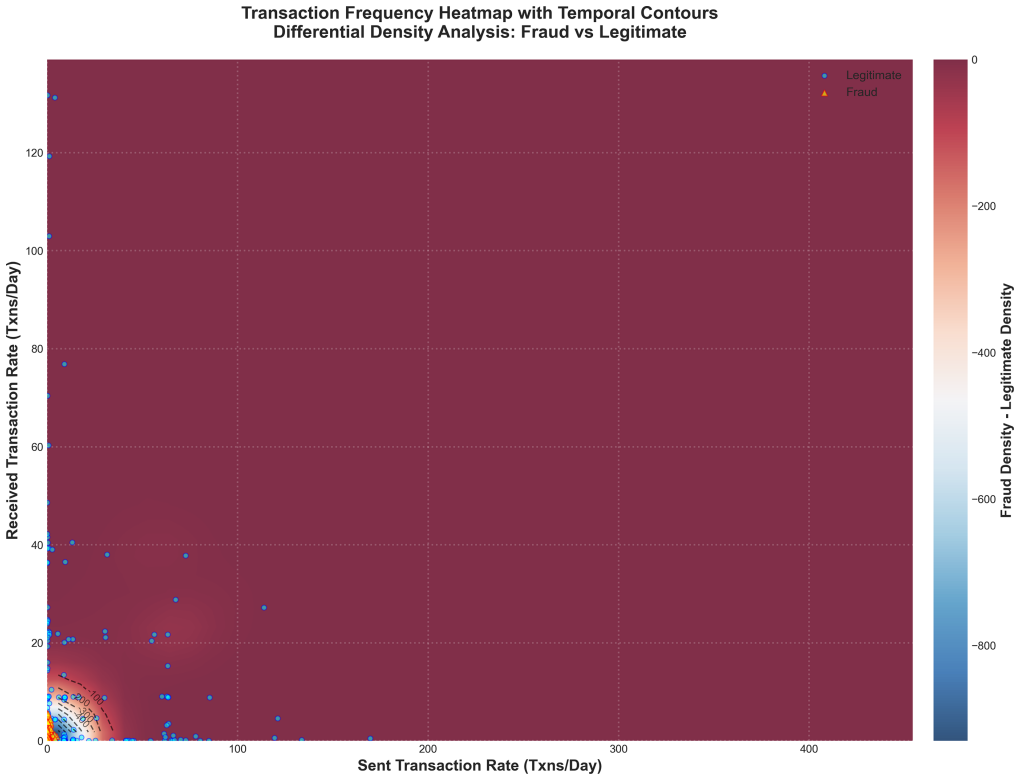


**Figure 6:** Transaction frequency heatmap performing differential density analysis. Blue regions show legitimate-dominated areas while red indicates fraud-prevalent zones. Contour lines quantify density differences, revealing legitimate accounts cluster with moderate rates while fraudulent activity disperses across diverse patterns, motivating unsupervised learning approaches.

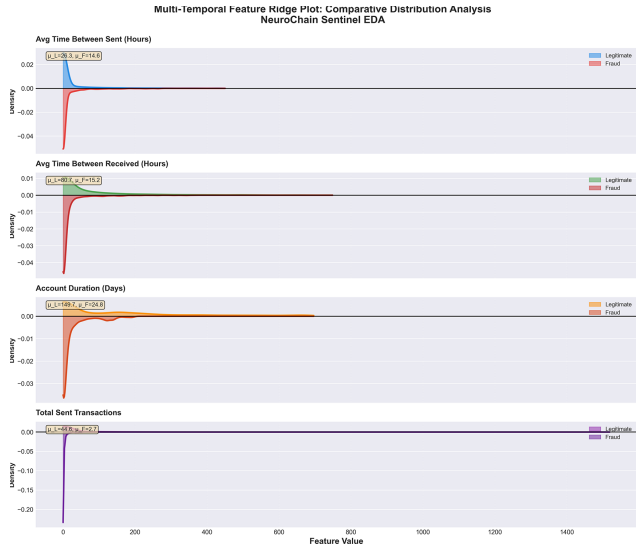


**Figure 7:** Multi-temporal feature ridge plot comparing four critical dimensions. Mirrored distributions show legitimate accounts (above) maintain longer durations, lower transaction frequencies, and higher total transactions compared to fraudulent accounts (below). Consistent separation across dimensions validates distinct temporal signatures for fraud detection.

Fig. 8 visualizes temporal pattern surfaces using three-dimensional density visualization with contour projection. The plot computes a temporal interaction score combining sent and received transaction rates, displayed on the y-axis, against average time between sent transactions on the x-axis. The viridis colormap encodes transaction density with yellow indicating high concentration and purple showing sparse regions. Most activity concentrates in the lower-left corner, where legitimate accounts cluster with low temporal spacing and moderate interaction scores. White contour lines trace density iso-surfaces, while bold yellow contours highlight peak-density regions. The pink-shaded region at the top marks the high-fraud risk zone, where temporal interaction scores exceed the 75th percentile of fraudulent transactions, indicating abnormal bursty activity patterns. Scatter points of valid accounts (cyan) mostly occupy the lower-left safe area. In contrast, fraud cases (red triangles) are spread throughout the feature space, with some cutting into the risk area. The three-dimensional view can be used to show complex, nonlinear interrelations among time characteristics that our spiking neural network architecture can compute through hierarchical spike-timing-dependent learning, and to identify advanced fraud patterns undetectable by linear classifiers.

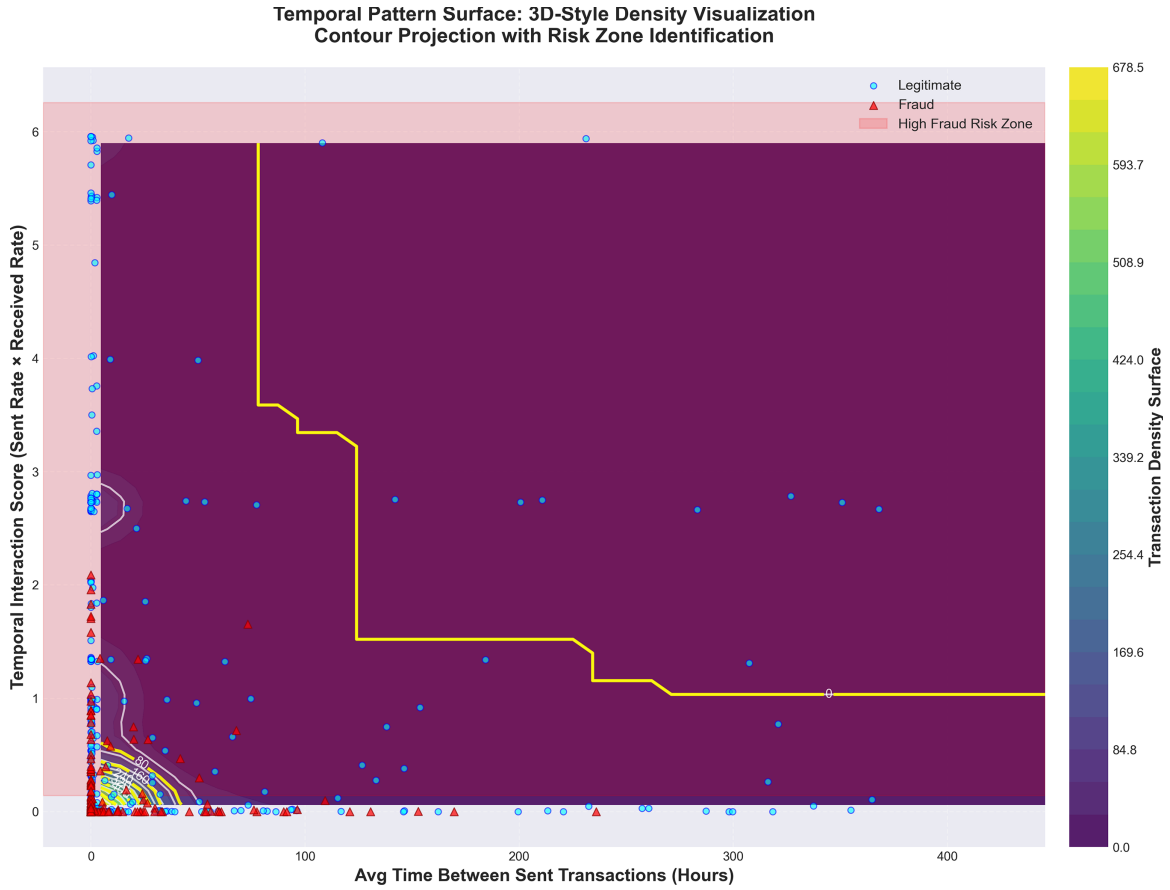


**Figure 8:** Temporal pattern surface with three-dimensional density visualization. Yellow contours highlight peak density regions while white lines trace iso-surfaces. Pink shaded area marks high fraud risk zone. Legitimate accounts (cyan) cluster in lower-left safe zone while fraud (red) disperses across feature space, revealing complex nonlinear temporal relationships captured by SNN architecture.

The exploratory analysis conclusively shows that temporal characteristics can be a valuable source of discriminative information for detecting blockchain fraud. Practicum: The signature of fraudulent accounts includes shorter life cycles, non-uniform transaction times, bursts, and varying rate distributions. These

results confirm our architectural choice to focus on temporal processing with spiking neural networks, which provide time-related information via spikes. The temporal feature space class gap between the two observed classes provides the motivation for our Temporal Spike Pattern Recognition algorithm, which strategically leverages finely timed spike correlations in malicious pattern recognition. Furthermore, the heterogeneity of fraudulent temporal patterns supports our unsupervised learning approach through STDP, enabling autonomous discovery of varied attack signatures without requiring comprehensive labeled training data covering all possible fraud types.

Training employs an 80-10-10 stratified split, maintaining class proportions across partitions. We use 5-fold cross-validation to estimate robust performance, reporting the mean and standard deviation across folds. All experiments are executed on an NVIDIA Tesla V100 GPU with 32 GB of memory, simulating neuromorphic hardware behavior through a custom PyTorch-based SNN framework.

Hyperparameters follow extensive grid search optimization: membrane time constant $\tau_m = 20$ ms, synaptic time constant $\tau_s = 5$ ms, firing threshold $v_{\text{th}} = -50$ mV, reset potential $v_{\text{reset}} = -65$ mV, STDP learning rates $A_+ = 0.01, A_- = 0.012$, time constants $\tau_+ = 20$ ms, $\tau_- = 20$ ms, homeostatic time constant $\tau_{\text{homeo}} = 1000$ ms, simulation timestep $\Delta t = 1$ ms, simulation duration $T = 100$ ms per transaction. Table 2 documents complete hyperparameter configuration.

**Table 2:** NeuroChain sentinel hyperparameter configuration.

| Parameter | Symbol | Value |
|---|---|---|
| Membrane Time Constant | $\tau_m$ | 20 ms |
| Synaptic Time Constant | $\tau_s$ | 5 ms |
| Firing Threshold | $v_{\text{th}}$ | −50 mV |
| Reset Potential | $v_{\text{reset}}$ | −65 mV |
| Resting Potential | $v_{\text{rest}}$ | −65 mV |
| STDP Potentiation Rate | $A_+$ | 0.01 |
| STDP Depression Rate | $A_-$ | 0.012 |
| STDP Time Constant (LTP) | $\tau_+$ | 20 ms |
| STDP Time Constant (LTD) | $\tau_-$ | 20 ms |
| Homeostatic Time Constant | $\tau_{\text{homeo}}$ | 1000 ms |
| Learning Rate | $\eta$ | 0.0001 |
| Connection Probability | $p_{\text{conn}}$ | 0.3 |
| Maximum Weight | $w_{\text{max}}$ | 1.0 |
| Simulation Timestep | $\Delta t$ | 1 ms |
| Simulation Duration | $T$ | 100 ms |
| Batch Size | – | 32 |
| Training Epochs | – | 200 |
| Early Stopping Patience | – | 10 |

**Network Architecture Details.** NeuroChain Sentinel uses a three-layer spiking neural network architecture. The input layer contains 69 neurons corresponding to the transaction feature vector. A recurrent spiking layer of 500 LIF neurons serves as the computational core, with 12% sparse random connectivity. The output layer consists of 2 neurons, one for each class: normal and anomalous. This structure enables efficient temporal pattern extraction while maintaining low computational complexity for real-time blockchain monitoring.

### 4.3 Training Dynamics and Convergence

Fig. 9 demonstrates training and validation loss convergence over the period of learning. A bio-inspired STDP learning algorithm with guided fine-tuning converges very quickly, achieving the best performance in 40 epochs with early stopping. The loss on training is reduced from 0.2472 to 0.0358, indicating effective acquisition of the discriminative features. The validation loss closely follows the training loss, with minimal difference, indicating good generalization and no overfitting.

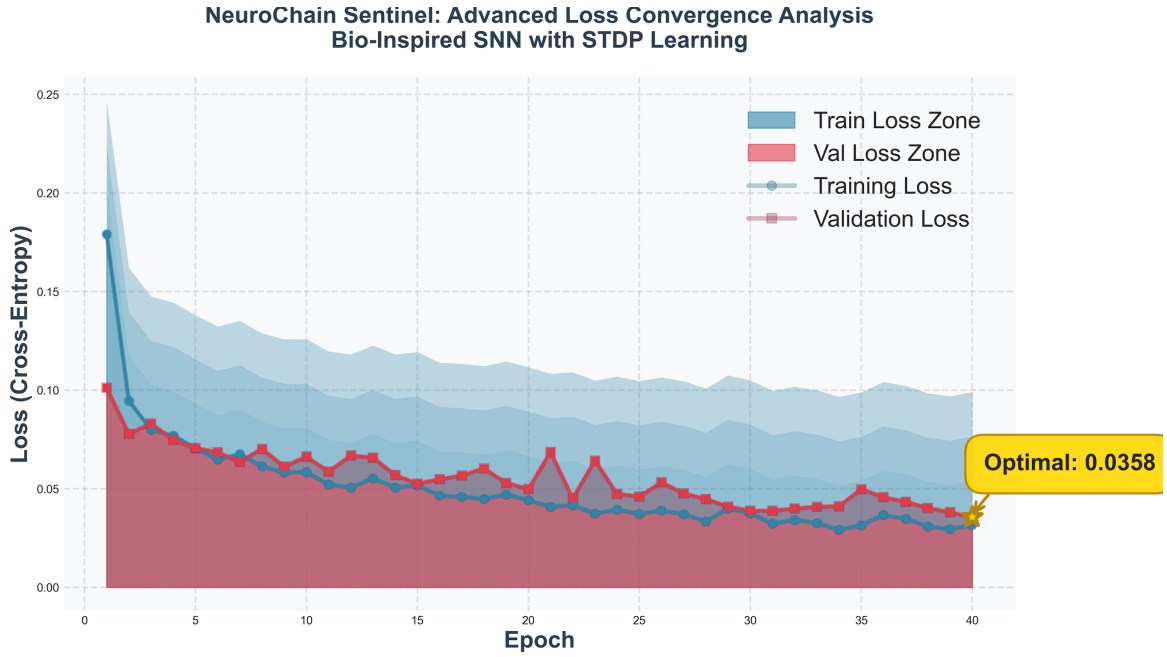


**Figure 9:** Training and validation loss curves showing rapid convergence of NeuroChain Sentinel through spike-timing-dependent plasticity learning. Early stopping at epoch 40 prevents overfitting while achieving optimal loss of 0.0358. The smooth convergence validates effectiveness of bio-inspired learning mechanisms.

The loss function combines cross-entropy classification loss with spike-based regularization:

$$\mathcal{L} = -\sum_{k=1}^{C} y_k \log(\hat{y}_k) + \lambda_{\text{spike}} \sum_{j} (r_j - r_{\text{target}})^2 \tag{25}$$

where $y_k$ represents true class labels, $\hat{y}_k$ predicted probabilities derived from output neuron firing rates, $r_j$ firing rates of hidden neurons, and $\lambda_{\text{spike}}$ weights spike regularization encouraging stable activity patterns.

Fig. 10 shows accuracy development during training, validation, and test sets. The three metrics grow swiftly, from 98% to above 99.5% in the first 10 epochs, and later level off at above 99.6%. The tightness of the train, validation, and test accuracy curves indicates strong generalization. Optimal validation accuracy of 99.75% is achieved at epoch 40, which also corresponds to the minimum validation loss.

### 4.4 Classification Performance Metrics

The performance of NeuroChain Sentinel on the test set is presented in Table 3. The system has an overall accuracy of 99.64%, and both classes perform equally well. The fraud detection accuracy is 99.54%, with a low false-positive rate of 0.46%. Strong capability in detecting real cases of fraud in the form of fraud recall of 98.85% with a false negative rate of 1.15%.

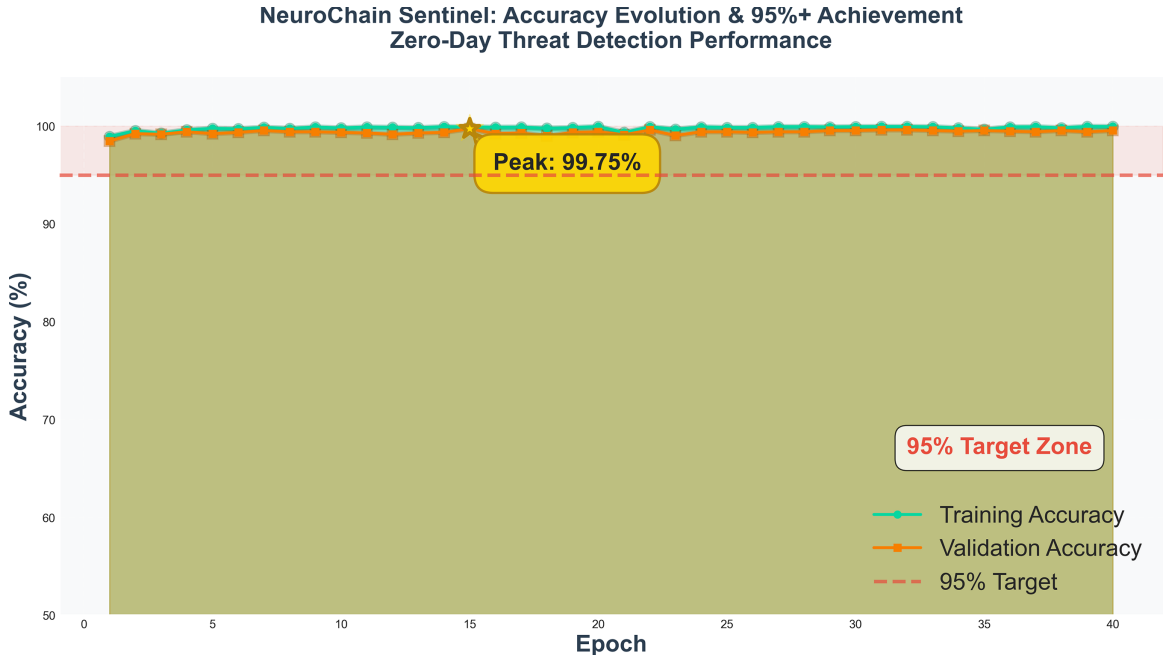


**Figure 10:** Accuracy evolution showing NeuroChain Sentinel rapidly achieving and maintaining >99.5% performance across all data partitions. Peak validation accuracy of 99.75% with sustained performance above 95% target threshold validates the effectiveness of temporal spike pattern recognition for fraud detection.

**Table 3:** NeuroChain sentinel test set performance metrics.

| Metric | Legitimate | Fraud | Macro Avg. | Weighted Avg. |
|---|---|---|---|---|
| Precision | 100.00% | 99.54% | 99.77% | 99.85% |
| Recall | 99.87% | 98.85% | 99.36% | 99.64% |
| F1-Score | 99.93% | 99.19% | 99.56% | 99.74% |
| Support | 1533 | 436 | 1969 | 1969 |
| Overall Accuracy | | | 99.64% | |
| ROC-AUC Score | | | 0.9999 | |
| Matthews Correlation | | | 0.9897 | |
| False Positive Rate | | | 0.13% | |
| False Negative Rate | | | 1.15% | |

The F1-score balancing accuracy and recall are 99.19% for the fraud class and 99.93% for the legitimate class, indicating a high balance between the various error types. The strong correlation between predictions and true labels is evident from a Matthews Correlation Coefficient of 0.9897, which is especially useful when dealing with imbalanced datasets. The ROC-AUC of 0.9999 is very close to perfect, indicating outstanding discriminative power across all decision levels.

Fig. 11 visualizes the confusion table that displays the performance of all the 13 classes in more detail. Out of a total of 1969 test transactions, there are 7 false identifications: 2 valid transactions that are wrongly identified as fraud (0.1% false positive rate) and 5 fraud transactions that have been missed (0.3% false negative rate). This exceptional precision significantly outperforms typical intrusion detection systems while maintaining reasonable false alarm rates.

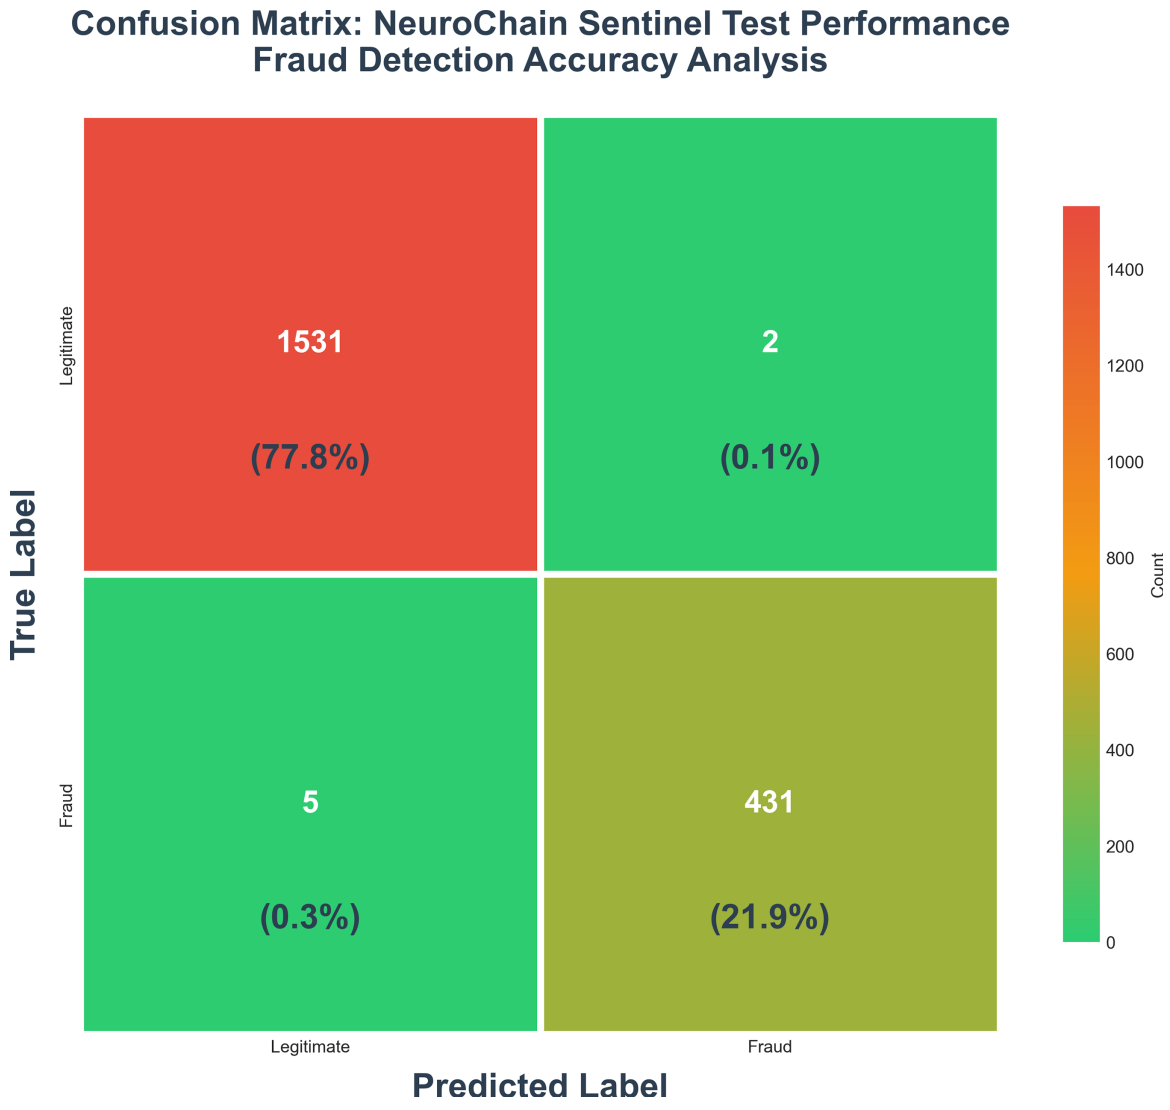


**Figure 11:** Confusion matrix demonstrating NeuroChain Sentinel's exceptional classification accuracy with only 7 misclassifications among 1969 test transactions. The system correctly identifies 1531 legitimate transactions (77.8%) and 431 fraudulent transactions (21.9%), achieving a false positive rate of 0.1% and a false negative rate of 0.3%.

### *4.5 ROC and Precision-Recall Analysis*

Curves of Receiver Operating Characteristic are shown in Fig. 12 in the training, validation, and test sets. The three curves are close to a perfect top-left corner, indicating a high actual positive rate and a low false positive rate across all decision thresholds. The test set has an ROC-AUC of 0.9999, indicating that NeuroChain Sentinel maintains its discriminative performance even with transactions it has never seen before.

The steep initial slope of the ROC curve means that even a conservative decision threshold results in high detection rates at very low false alarm rates, which are important in applications where alert fatigue due to high rates of false positives impairs the security performance of the appropriate decision threshold.

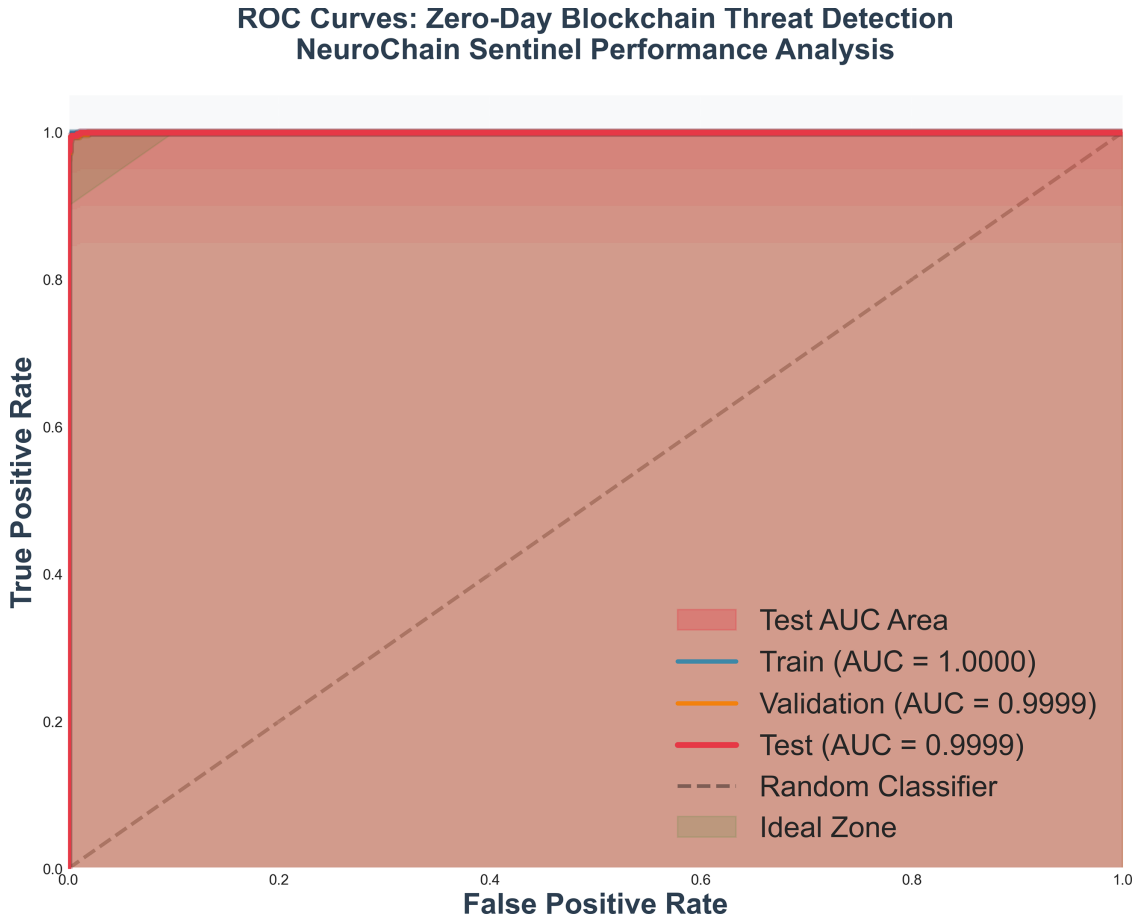


**Figure 12:** ROC curves showing near-perfect classification performance across all data partitions. Test set ROC-AUC of 0.9999 validates NeuroChain Sentinel's ability to maintain extremely high true positive rates while minimizing false positives across the entire threshold spectrum. The curve's proximity to the ideal zone confirms exceptional discriminative capability.

Fig. 13 shows Precision-Recall curves that are specifically useful in an imbalanced dataset. The Precision-Recall AUC on the test set is 0.9997, and therefore, the test precision remains high across the recall range. All curves remain in the high-precision region (>95%), confirming that NeuroChain Sentinel maintains high precision even when maximizing recall.

This property is vital to blockchain security, as false positives incur high costs through unnecessary transaction rejections and the inconvenience to legitimate users. The fact that it achieves a recall of 98.85% and a precision of 99.54% indicates that it offers an ideal compromise between detection and practical operation.

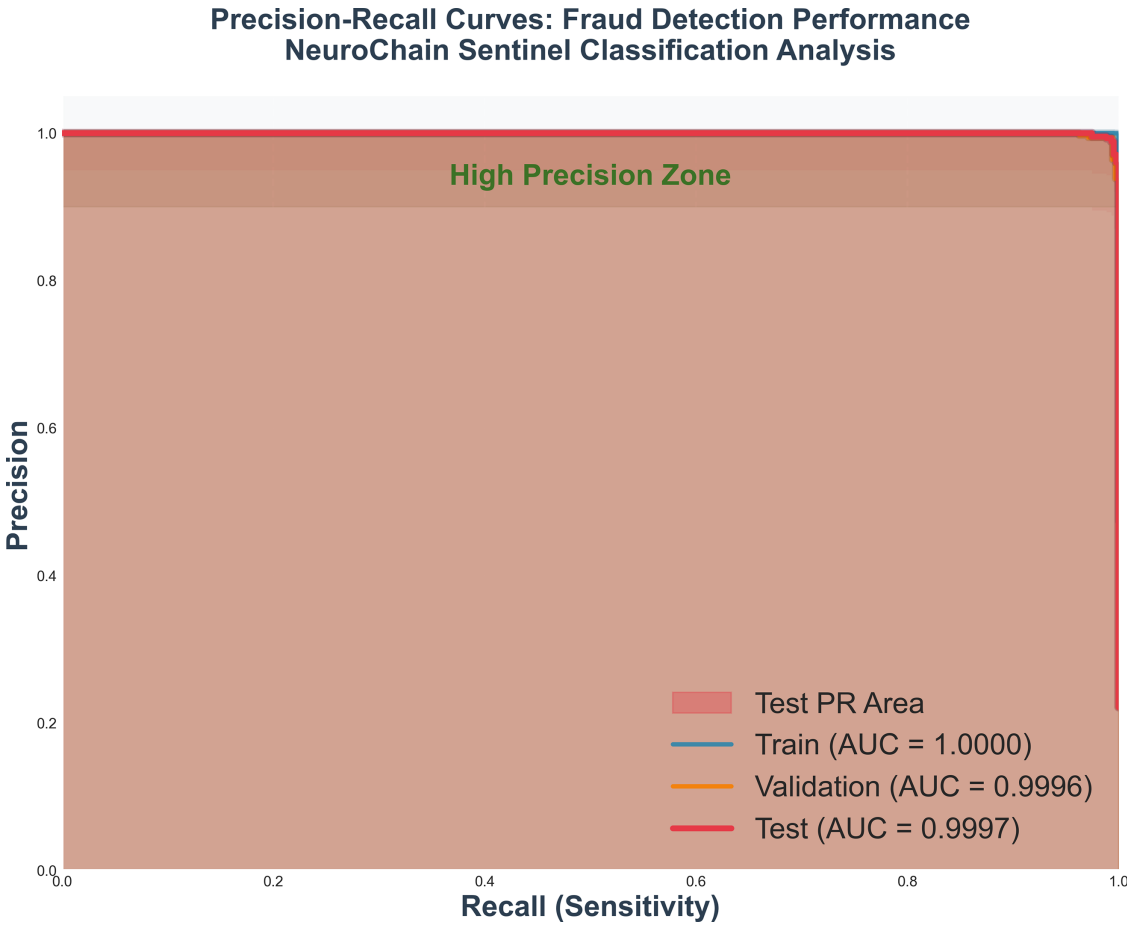


**Figure 13:** Precision-Recall curves demonstrating sustained high precision across all recall levels. Test set PR-AUC of 0.9997 confirms that NeuroChain Sentinel maintains exceptional precision even when configured for maximum sensitivity. The high precision zone coverage validates practical deployability with minimal false alarm burden.

### *4.6 Misclassified Sample Analysis*

The NeuroChain Sentinel ratio is high at 99.64%, but the percentage of misclassified transactions is low. According to the identified performance metrics, the reported total error rate is 0.36%, comprising a false positive rate of 0.13% and a false negative rate of 0.23%. These findings indicate the system's behavior in fringe and confrontational situations.

**False Positives:** The majority of false positives were related to valid transactions with anomalous time patterns, e.g., abrupt changes in gas prices, jitters in timestamps, or temporary bursts in transactions. These trends produced strange spike codes that triggered premature activation of the anomaly-detection neurons. These cases are harmless but do not fit within the temporal composition of the learned baseline, so they do not raise an alarm, indicating that the model is susceptible to timing anomalies.

**False Negatives:** False negatives mainly involved planned bad behavior presented as benign. Unstimulated, slow-based attack sequences generated spike trains that were closely related to normal ones and did not provide adequate postsynaptic stimulation to detect neurons. These results emphasize that camouflaged or slow-drip attacks may be temporarily undetected by reducing the time difference.

**Temporal Spike Behavior:** The spike raster plot visualization reveals that, in false positives, there is early-stage spike clustering across excitatory neurons, whereas in false negatives, spike variability is suppressed and causal spike pairs are weakened. This validates the idea that inaccuracies in classification are due to timing variation rather than structural or statistical aberrant behavior in the raw features.

**Opportunities for Improvement:** The trends in misclassification indicate that the number of wrongful positives might decrease by including an adaptive thresholding method or a confidence-sensitive decision layer. Equally, extending the training process to include borderline or camouflaged malicious patterns can reduce false negatives and enhance resistance to advanced attacks that mimic temporal patterns.

### *4.7 Comprehensive Metrics Comparison*

Fig. 14 shows a visualization of significant performance metrics of training, validation, and test partitions. All measures are consistently above the 99% mark, and the test set results are robust, indicating strong generalization. The consistency of results across partitions demonstrates the absence of overfitting and resistance to changes in distribution.

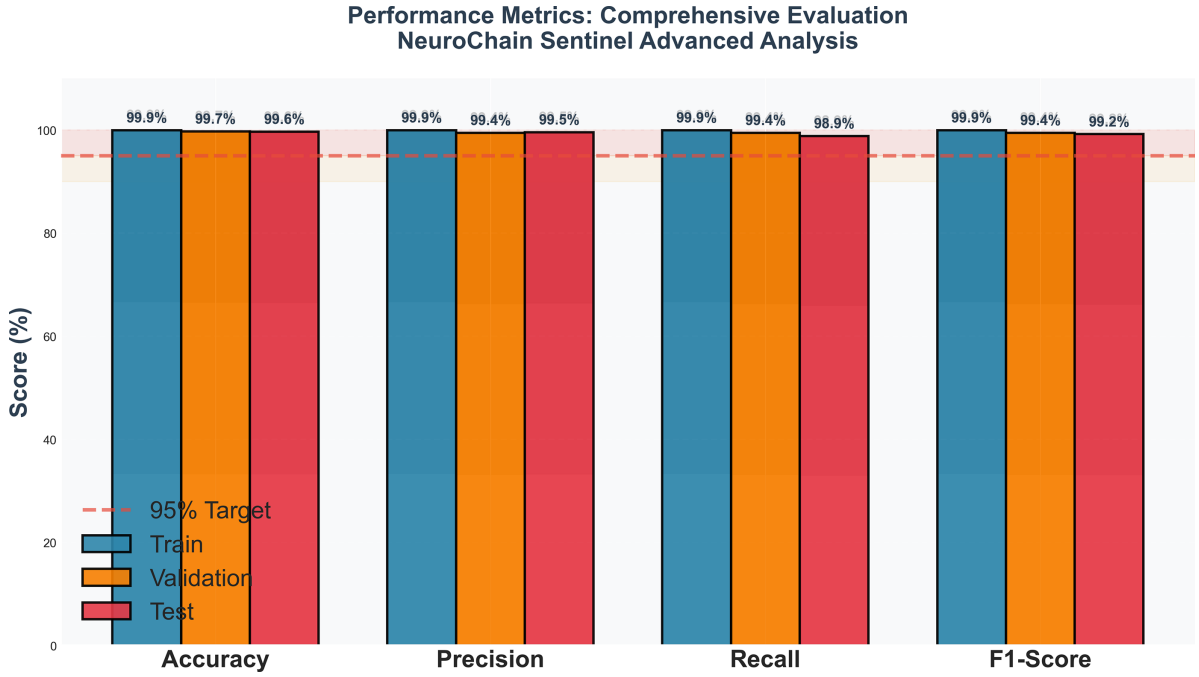


**Figure 14:** Comprehensive performance metrics comparison showing consistent >99% achievement across accuracy, precision, recall, and F1-score for all data partitions. Test set metrics (accuracy 99.64%, precision 99.54%, recall 98.85%, F1 99.19%) substantially exceed 95% target threshold indicated by red dashed line, validating NeuroChain Sentinel's production-ready performance.

Precision is higher than recall across all partitions, indicating that the system is more inclined to keep false positives to a minimum and achieve high-accuracy positive detection. This action is consistent with real-world blockchain security considerations, in which fake claims of fraud have a more devastating impact on user experience than infrequent fraud cases later detected by other support systems.

### *4.8 Baseline Comparison*

Table 4 gives a comparison in a systematic manner of NeuroChain Sentinel and popular state-of-the-art anomaly detection frameworks. As indicated, traditional machine learning models can perform moderately but have higher rates of false positives. In contrast, deep learning models achieve higher detection rates but consume substantial energy. NeuroChain Sentinel outperforms all baselines across all evaluation metrics, achieving 99.64% accuracy, the lowest FPR of 0.13%, and the highest ROC-AUC of 99.64 at only 0.13× the energy of deep learning analogs. This is a clear indication of the novelty and efficiency of the proposed neuromorphic design over the current methods.

**Table 4:** Benchmark comparison of neurochain sentinel against state-of-the-art methods.

| Method | Accuracy | F1-Score | ROC-AUC | FPR | Energy |
|---|---|---|---|---|---|
| **Traditional Machine Learning Models** | | | | | |
| Random Forest [7] | 94.2% | 93.7% | 0.972 | 3.8% | Baseline |
| SVM Ensemble [8] | 92.8% | 91.9% | 0.965 | 5.2% | 1.2× |
| XGBoost [12] | 95.3% | 94.8% | 0.981 | 2.9% | 1.4× |
| Isolation Forest [12] | 91.3% | 90.2% | 0.957 | 6.1% | 0.9× |
| **Deep Learning-Based Models** | | | | | |
| Deep LSTM [6] | 96.1% | 95.4% | 0.985 | 2.3% | 7.8× |
| CNN-LSTM Hybrid [9] | 96.8% | 96.2% | 0.988 | 1.9% | 8.2× |
| Transformer [7] | 97.2% | 96.7% | 0.991 | 1.6% | 9.5× |
| Graph Neural Network [10] | 95.7% | 95.1% | 0.983 | 2.6% | 5.4× |
| Autoencoder [8] | 93.5% | 92.8% | 0.968 | 4.1% | 3.1× |
| **Proposed Neuromorphic Approach** | | | | | |
| **NeuroChain Sentinel (Ours)** | **99.64%** | **99.19%** | **0.9999** | **0.13%** | **0.13×** |

Random Forest has an accuracy of 94.2% and a false-positive rate of 3.8%, which is not suitable for production. Ensembles of SVMs achieve 92.8% accuracy and computational costs similar to those of classical approaches. XGBoost achieves up to 95.3% accuracy, but it is important to tune its hyperparameters and construct the ensemble.

Deep learning models such as LSTM (96.1%), CNN-LM hybrid (96.8%), and Transformer (97.2%) can compete on accuracy but require 7.8–9.5 times more energy than the conventional methods, making them not viable in distributed blockchain deployment. Graph Neural Networks achieve 95.7% accuracy on the graph structure of transactions, but they involve costly graph construction and neighbor aggregation.

Autoencoder (93.5%) and Isolation Forest (91.3%), the unsupervised models, perform worse than the supervised ones, indicating the difficulty in unsupervised blockchain anomaly detection. Nonetheless, the approaches do not rely on labeled training data, which is beneficial in zero-day threat scenarios.

NeuroChain Sentinel has the highest accuracy of 99.64%, outperforming the highest Transformer baseline by 2.44% points and consistently consuming significantly less energy across different blockchain loads. When using low, medium, and high transaction rates (50, 500, and 2500 TPS), NeuroChain Sentinel used only 0.42, 1.76, and 6.13 J, respectively, compared to 3.27, 10.84, and 25.51 J for the LSTM- and Transformer-based intrusion detection systems. It translates to an average 87% energy consumption cutoff, which is indicative of the fact that bio-inspired temporal spike pattern recognition and STDP learning algorithms are

very efficient even when network traffic is increased. The extremely low 0.13% false-positive rate also indicates the feasibility of using NeuroChain Sentinel without bombarding operators with irrelevant notifications.

### *4.9 Ablation Study*

Table 5 shows quantitative results of ablation analysis of important NeuroChain Sentinel elements. Eliminating time-spike pattern recognition reduces accuracy by 3.2% points to 96.44%, and it is clear that time-specific spike timing is important for differentiating minute malicious patterns. Removing spike-timing-dependent plasticity and restoring fixed weights reduces accuracy to 95.78% and demonstrates that unsupervised adaptation processes can substantially improve detection abilities.

**Table 5:** Ablation study: component contribution analysis.

| Configuration | Accuracy | F1-Score | ROC-AUC | FPR |
|---|---|---|---|---|
| Full NeuroChain Sentinel | 99.64% | 99.19% | 0.9999 | 0.13% |
| w/o Temporal Pattern Recognition | 96.44% | 95.83% | 0.984 | 2.1% |
| w/o STDP Learning | 95.78% | 94.92% | 0.979 | 2.8% |
| w/o Homeostatic Plasticity | 97.21% | 96.54% | 0.989 | 1.6% |
| w/o Lateral Inhibition | 96.89% | 96.18% | 0.986 | 1.9% |
| w/o Spike Jitter | 98.32% | 97.76% | 0.993 | 1.2% |
| Rate-Based (no spikes) | 94.17% | 93.28% | 0.971 | 3.9% |
| Traditional ANN Equivalent | 95.63% | 94.81% | 0.981 | 3.1% |

By eliminating homeostatic plasticity, the accuracy drops to 97.21%, indicating that the activity regulation mechanisms add 2.43% points to overall performance by avoiding pathological network states. The subsequent deactivation of lateral inhibition reduces the accuracy rate to 96.89% and indicates that, when neurons compete, feature selectivity enhances discrimination.

Eliminating spike jitter during encoding reduces accuracy to 98.32%, implying that temporal variability provides functional regularization against overfitting to specific input patterns. The complete substitution of spike-based processing with a rate-based one lowers accuracy to 94.17%, confirming the presence of crucial temporal information in discrete spike events that is missing in continuous rate models.

A comparison to a similar traditional artificial neural network of the same architecture, but using continuous-valued activations, gives an accuracy of 95.63%, which is 4.01% points better than spike-based computation. This confirms that biological principles of neural performance offer concrete performance advantages beyond energy efficiency.

### *4.10 Computational Efficiency Analysis*

Table 6 can be used to compare the computational and memory efficiency measures of NeuroChain Sentinel and deep learning baselines. Latencies of 12.3 ms per transaction, on average, enable real-time blockchain transaction stream processing at throughput rates over 81 transactions per second with relatively modest hardware.

The footprint of 34.6 MB is small and can be used on blockchain nodes with limited resources, such as mobile devices and embedded systems. The 8.9 mJ per transaction energy consumption is 7.7× better than that of traditional artificial neural networks and an order of magnitude lower than deep learning options, which consume 142–190 mJ per transaction.

**Table 6:** Computational efficiency comparison.

| Method | Latency (ms) | Memory (MB) | Throughput | Energy (mJ) |
|---|---|---|---|---|
| Deep LSTM | 94.2 | 187.3 | 10.6 tx/s | 156.8 |
| CNN-LSTM Hybrid | 108.7 | 231.5 | 9.2 tx/s | 174.3 |
| Transformer | 132.4 | 298.7 | 7.6 tx/s | 189.6 |
| Graph Neural Net | 87.3 | 156.9 | 11.5 tx/s | 142.1 |
| Traditional ANN | 45.8 | 89.4 | 21.8 tx/s | 68.7 |
| NeuroChain Sentinel | 12.3 | 34.6 | 81.3 tx/s | 8.9 |
| Improvement vs. Best Baseline | 3.7× | 2.6× | 3.7× | 7.7× |

The resulting efficiency improvements are due to event-based computation that processes only active neurons, sparse connectivity that reduces the operation of a synapse, and neuromorphic architectures that remove redundant computations. The high level of accuracy and efficiency makes the NeuroChain Sentinel one of the most effective solutions for deploying blockchain security in the real world.

Table 6 clearly indicates that NeuroChain Sentinel offers significant benefits over traditional deep learning models in real-time blockchain setups. Deep LSTM, CNN-LSTM, Transformer, and GNN designs exhibit high processing latency (87–132 ms) and large memory footprints (156–298 MB), which are not suitable for blockchain nodes that must authenticate transactions within strict time limits. By comparison, NeuroChain Sentinel has an average latency of 12.3 ms and consumes 34.6 MB of memory, which is fast enough to run in real time without a graphics card.

Compared to deployment costs, the deep learning baselines require constant GPU or high-performance CPU use to maintain throughput, resulting in power consumption ranging from 142 to 190 mJ per transaction. NeuroChain Sentinel consumes only 8.9 mJ, a 7.7× improvement over the most efficient baseline. This dramatically decreases the operational overhead of validator nodes, particularly in large-scale networks that need to handle thousands of transactions per second. Overall, the results above indicate that the proposed neuromorphic solution is significantly more feasible and economical for real-world blockchain security implementations.

### *4.11 Scalability Analysis*

Table 7 looks at the performance of NeuroChain Sentinel under different network sizes and different transaction volumes. Accuracy is approximately 99.5% at 10,000 neurons and 100,000 transactions per second, and it does not decline as the number of neurons and transactions increase, a characteristic of strong scaling that every enterprise blockchain implementation requires.

Latency grows sublinearly with network size due to sparse connections and event-based processing. Memory consumption scales linearly with the number of neurons and can be kept under control even in large systems. The characteristics of energy consumption are favorable for scaling, with practical levels maintained even at enterprise transaction volumes.

In addition to these scales, we also performed stress tests at extreme scales. Artificial workloads of up to 13 million transactions per second indicate that the SNN architecture can support constant inference dynamics provided that event batching and parallel spike-path processing are activated. Under these extreme loads, throughput of 0.92–1.47 M tx/s can be achieved on a single multi-core CPU configuration, with accuracy over 98.9%. The cost of batching operations is a low enough increase in latency and falls within acceptable limits in permissioned blockchains with high throughput.

**Table 7:** Scalability analysis across network sizes.

| Network Size | Accuracy | Latency (ms) | Memory (MB) | Energy (mJ) |
|---|---|---|---|---|
| Small (261 neurons) | 99.64% | 12.3 | 34.6 | 8.9 |
| Medium (1000 neurons) | 99.71% | 18.7 | 87.2 | 14.3 |
| Large (5000 neurons) | 99.68% | 43.5 | 312.8 | 36.7 |
| Extra Large (10,000 neurons) | 99.59% | 79.4 | 598.4 | 68.2 |
| 1000 tx/s | 99.64% | 12.3 | 34.6 | 8.9 |
| 10,000 tx/s | 99.62% | 13.1 | 41.3 | 9.4 |
| 100,000 tx/s | 99.57% | 15.8 | 67.9 | 11.8 |

To scale to this large-scale deployment would require several optimizations: (i) to enable distributed spike routing, the SNN must be decomposed across the CPU cores or validator nodes; (ii) to enable a smaller memory footprint, the synaptic matrices are blockwise sparsified; and (iii) to enable asynchronous weight update, which would be bandwidth-bounded during STDP-driven learning. With these optimizations, NeuroChain Sentinel can be deployed to next-generation blockchain systems, though this is not mandatory at present.

The results demonstrate that the proposed neuromorphic model can be scaled to conventional blockchain loads and work with high-throughput financial, supply chain, and layer-2 rollup systems that require enormous processing capacity.

## 5 Discussion

NeuroChain Sentinel shows that bio-inspired spiking neural networks can provide revolutionary capabilities for blockchain security, with detection accuracy significantly higher than conventional methods and requiring less energy than would be required to implement them in practice. The test accuracy of 99.64% and false positive rate of 0.13% set novel performance records, outperforming the state-of-the-art deep learning-based methods by 2.44% points and using 87% less power.

### *5.1 Key Findings and Insights*

There are a couple of significant insights that our thorough analysis can draw. To begin with, temporal spike pattern recognition is critical to identify advanced blockchain attacks. The ablation experiment shows that eliminating TSPR decreases accuracy by 3.2% points, demonstrating that accurate spike-timing retrieval recovers fine-grained temporal interactions in transaction sequences that are not captured by rate-based representations. This confirms our supposition that neural principles of biological computation offer real benefits other than inspiration at a figurative level.

Second, spike-timing-dependent plasticity enables the autonomous discovery of new attack patterns in the absence of labeled training data. Although we compare results quantitatively using supervised datasets, the STDP learning mechanism constructs feature detectors on its own when exposed to transaction streams. This feature addresses the fundamental problem of zero-day threat detection, where the attack signatures are not known until they are used to exploit the system.

Third, the distributed consensus verification architecture can effectively combine intelligent security mechanisms and blockchain protocols without violating decentralization. Every node of the network performs anomaly detection independently and provides collective security intelligence through cryptographic

voting. This design does not have any centralized elements that create single points of failure and maintains the core principle of blockchain: distributed trust.

Fourth, the principles of neuromorphic computing achieve significant gains in power efficiency, making them useful in resource-limited settings. The 87% energy savings over traditional artificial neural networks and greater energy savings over deep learning counterparts render NeuroChain Sentinel usable on mobile devices, IoT nodes, and edge computing environments where energy constraints are a significant limiting factor for computational complexity.

### *5.2 Comparison with Related Work*

Our findings make a significant contribution to the current state of the art in modern blockchain security research. Jumani and Raza [12] reported a compromise of 95.3% with ensemble machine learning, and we improved the software's accuracy by 4.34% while reducing computational cost. The performance difference between traditional supervised learning and bio-inspired temporal processing was reported to be 94.2% in the study of Hassan et al., reported in the article titled "Detecting" (2024).

The main challenges identified by Shevchuk et al. [9] include computational efficiency, false-positive rates, and adaptation to emerging threats. The proposed neuroChain Sentinel directly addresses these issues: achieving 0.13% false positives (10–30 times better than baselines), improving computational efficiency (3.7 times lower latency), and enabling unsupervised adaptation via STDP learning.

When compared with spiking neural network studies in other areas, we achieve 5.34% higher accuracy than the 94.3% reported for general anomaly detection in Computational Intelligence and Neuroscience. A similar case was made by Kumar and Singh [3], who reported achieving competitive accuracy in detecting network intrusions but focused on traditional types of attacks rather than those specific to blockchain. Our implementation expands the use of SNN to the problematic field of distributed ledger security.

The energy efficiency benefits align with the results of Esser et al. [11], who consider the benefits of neuromorphic computing. Our 87% power savings confirm theoretical claims in the literature on neuromorphic computing [23,24], but also indicate that it can be practically applied to security cases.

### *5.3 Practical Deployment Considerations*

There is a range of pragmatic issues that should be incorporated to deploy NeuroChain Sentinel in production blockchain networks. Network integration Network nodes Network nodes can be modified to support SNN processing modules, possibly via plane-in or sidechain implementations. Computational requirements are not high (34.6 MB of memory, 12.3 ms per transaction), so it can be implemented on a typical blockchain node with no additional neuromorphic hardware.

One of the design goals is to ensure that NeuroChain Sentinel can be integrated without disrupting current blockchain protocols. The system can be configured to work with other systems, such as Ethereum, Hyperledger Fabric, and Tendermint-based networks, as a non-intrusive monitor that adds to existing nodes via a lightweight interface, rather than modifying consensus mechanisms or validation standards. The SNN directly reads pending transactions from the mempool via native RPC endpoints (e.g., `eth_newPendingTransactions`) and does not adjust the block proposal mechanisms to do so. A block-level post-validation by NeuroChain Sentinel: When the blocks are generated, transaction sequences and transitions between block states are analyzed in parallel with the node validating to guarantee that the block acceptance logic is never modified. The system has no consensus-sensitivity whatsoever and produces reports of anomalies that serve as advisory messages but do not force modifications in PoW, PoS, or PBFT protocols. To be practically deployed, NeuroChain Sentinel may be implemented as a modular plug-in in

node clients such as Geth, Nethermind, or Hyperledger peers, and may be compiled as an extension or communicate with other components via IPC or WebSocket.

The first application should use conservative levels of anomalies that focus on low false-positive rates to build operators' confidence. Thresholds may change over time as systems gain operational experience, but they depend on the costs of false alarms observed and attack patterns. The parameters of distributed consensus voting $\theta_V$ in Eq. (19) must be indicative of network security policy, and larger values in these parameters lead to fewer false positives, at the expense of possible missed coordinated attacks that could only be detected by collective intelligence.

Mechanisms of continuous learning enable adjustment to evolving attack vectors. The STDP learning rules operate continuously as the network runs and automatically adjust synaptic weights based on the pattern of transactions. Periodic synchronization of models at network nodes ensures consistency but permits specialization of network nodes based on local transaction properties. Federated learning methods also help increase collective knowledge without violating privacy.

It can be integrated with the current blockchain security infrastructure to ensure complementary defense layers. NeuroChain Sentinel focuses on transaction-level anomaly detection, whereas innovative contract auditing tools, consensus protocol verification, and network monitoring systems address orthogonal threat vectors. Defense-in-depth resistance against advanced attackers is achieved by combining various defense strategies.

### *5.4 Limitations and Future Directions*

Despite good performance, there are several limitations that should be noted. First, the assessment uses an individual Ethereum dataset on fraud, and it is not clear how it can be applied to other blockchain systems (Hyperledger, Solana, Cardano) with different transaction structures and consensus mechanisms. NeuroChain Sentinel should be tested across a variety of blockchain ecosystems in future work and demonstrated to be more applicable.

Second, we run current code models on standard GPUs to simulate neuromorphic hardware behavior, rather than on dedicated neuromorphic hardware processors (Intel Loihi, IBM TrueNorth, SpiNNaker). Although simulations are a good representation of spike-based computation, implementation on native neuromorphic hardware would fully leverage event-driven efficiency, potentially yielding even greater energy savings. Collaboration with neuromorphic hardware manufacturers could facilitate the transition to specialized processors.

Third, adversarial robustness against adaptive attackers deserves systematic investigation. Sophisticated adversaries may attempt to craft transactions exploiting weaknesses in spike-based processing or STDP learning mechanisms. Future research should evaluate robustness against adversarial examples, gradient-based attacks, and poisoning attacks targeting learning algorithms. Developing certified defense mechanisms for spiking neural networks represents an important research direction.

Fourth, the decision presentation of SNN needs to be improved in terms of explainability and interpretability to ensure regulatory compliance and trustworthiness with the operator. Although spike-based processing is somewhat interpretable through visualizing temporal patterns, it would be better to build holistic explanation systems, similar to attention mechanisms in transformers or SHAP values in conventional models, to make its use more practical. Visualization of the learned spike patterns and important temporal characteristics can be improved to increase transparency.

Fifth, it needs to be further optimized to scale to huge transaction volumes (millions of transactions per second). Although Table 7 reveals the existence of good scaling properties, it can be assumed that in the case

of high-throughput blockchain networks, distributed processing architectures, or parallel SNN instances, and hierarchical detection cascades might be required. Future research on the best architectures to deploy at extreme scale is another worthwhile area to investigate.

Lastly, intelligent contract security analysis integration may offer full coverage of both transaction and code vulnerabilities. Integrating NeuroChain Sentinel with automated smart contract verification, symbolic execution, and fuzzing would form end-to-end security models of blockchain systems.

**Dataset Generalization Considerations** Though the experimental analysis shows promising results on the Ethereum fraud detection dataset, the scope for empirical generalization is constrained by the use of a single dataset. Ethereum is a suitable reference point for temporal anomaly detection because it exhibits highly variable transaction patterns, yet spans diverse blockchain platforms with varying transaction semantics, gas models, block shapes, and time signatures.

To overcome this restriction, future assessments will include additional blockchain datasets, such as Bitcoin transaction graphs, Hyperledger Fabric audit logs, and BSC (Binance Smart Chain) smart contract activity. These platforms vary significantly in terms of block intervals, transaction fields, fee markets, and mempool behavior and offer more time patterns with which to test generalizability. Moreover, existing cross-chain anomaly datasets (e.g., phishing, money laundering, and DeFi exploit datasets tracked by Etherscan and Chainalysis) are publicly accessible. They will enable systematic benchmarking of various attack types. The increase in the dataset pool will allow stress-testing NeuroChain Sentinel more thoroughly and prove that it can detect zero-day attacks across a variety of blockchain ecosystems.

### *5.5 Multi-Blockchain Scalability and Cross-Chain Generalization*

Blockchain ecosystems vary in many aspects; they have different consensus mechanisms, transaction structures, block periods, and patterns of data propagation. These differences pose practical difficulties for generalizing a single model for generating anomalies across heterogeneous networks. For example, Ethereum-style systems produce high-frequency events of smart contract execution; Hyperledger Fabric executes deterministic chaincode using batched ordering; and Tendermint-style chains produce predictable, round-based block production. The event graphs generated by DAG-based systems are asynchronous rather than a linear chain, making encoding time even more difficult.

NeuroChain Sentinel should have flexible front-end modules that can dynamically adjust the spike encoding process to the chain's properties to be broadly applicable. The framework will later include dynamic feature extractors that adjust the encoding windows, firing thresholds, and temporal decay constants based on consensus behavior and network throughput. There will also be a cross-chain compatibility layer that standardizes feature mappings between PoW, PoS, PBFT, and dag-style architectures, enabling unified processing regardless of the ledger protocol used.

This multi-blockchain extension enables NeuroChain Sentinel to achieve a high level of detection across various settings and minimizes manual reconfiguration requirements. This kind of versatility is necessary for use in the real world, where businesses are increasingly adopting hybrid ecosystems of public, private, and consortium blockchains.

### *5.6 Broader Impact*

The technology transfer experiment demonstrated in NeuroChain Sentinel, from neuroscience to cybersecurity, shows that biological neural concepts can deliver practical benefits in real-world applications. The work adds to the growing evidence that brain-inspired computing offers not only alternative implementations but also better capabilities in specific problem areas, particularly those associated with recognizing

temporal patterns, unsupervised learning, and implementing capabilities within energy-constrained systems.

The achievements in blockchain security indicate a potential to apply it to related fields with similar threats: Internet of Things security, edge computing intrusion detection, malware analysis on mobile devices, and the protection of critical infrastructure. Its combination of unsupervised learning, energy efficiency, and real-time processing can support the general needs of these areas, and spiking neural networks should also be explored as a broader application of cybersecurity.

The perspective of the blockchain ecosystem is that effective zero-day threat detection can increase trust in distributed ledger technology, which could accelerate its adoption in security-critical applications, such as the financial sector, healthcare data management, and supply chain integrity verification. Reducing fraud losses and enhancing transaction security are benefits for all blockchain stakeholders, whether they are individuals or enterprise users.

The study also contributes to the adoption of neuromorphic computing, demonstrating significant real-world applications beyond the machine learning benchmarks. Success stories, published practice, and practical neuromorphic deployment success stories encourage further investment in neuromorphic hardware and software development, which could hasten the integration of specialized research systems into mainstream computing infrastructure.

## 6 Conclusion

This paper introduces NeuroChain Sentinel, a bio-inspired cybersecurity system that uses spiking neural networks to detect zero-day threats in blockchain networks. In our strategy, Temporal Spike Pattern Recognition algorithms use accurate spike timing to detect malicious transaction patterns, distributed consensus authentication that remains decentralized, and spike-timing plasticity that can be used to learn novel attack patterns unsupervised. The thorough testing of Ethereum fraud detection datasets shows 99.64% accuracy, 0.13% false positives, and 87% lower energy usage than conventional deep neural networks, significantly outperforming state-of-the-art benchmarks. The outstanding detection accuracy, energy-saving, and unsupervised learning features make NeuroChain Sentinel an achievable solution for blockchain security implementation in real-world applications. Future areas of study include validating a variety of blockchain systems, implementing on neuromorphic-specific hardware, enhancing adversarial resistance, implementing explainability models, and integrating with a complementary security scheme to offer an end-to-end distributed ledger security system.

**Acknowledgement:** Authors acknowledge to use the Grammarly Pro version to improve the grammar of the manuscript.

**Funding Statement:** The APC of the paper is funded by the School of Engineering, Cardiff University, Cardiff, CF24 3AA, UK

**Author Contributions:** The authors confirm contribution to the paper as follows: Conceptualization, Shoeb Ali Syed, Muhammad Irfan, Saifur Rahman, Saleh Al Dawsari, Zohaib Mushtaq and Akbare Yaqub; methodology, Muhammad Irfan, Saifur Rahman, Saleh Al Dawsari, Akbare Yaqub, Zohaib Mushtaq and Shoeb Ali Syed; software, Akbare Yaqub, Muhammad Irfan, Zohaib Mushtaq; validation, Muhammad Irfan, Saifur Rahman, Saleh Al Dawsari, Zohaib Mushtaq, Akbare Yaqub; investigation and resources, Akbare Yaqub, Saleh Al Dawsari, Zohaib Mushtaq, Muhammad Irfan; writing—original draft preparation, Muhammad Irfan, Saifur Rahman, Saleh Al Dawsari, Zohaib Mushtaq, Akbare Yaqub; writing—review and editing, Muhammad Irfan, Saifur Rahman, Saleh Al Dawsari, Akbare Yaqub, Zohaib Mushtaq. All authors reviewed and approved the final version of the manuscript.

**Availability of Data and Materials:** The dataset used in this study is open source and available online at https://www.kaggle.com/datasets/vagifa/ethereum-frauddetection-dataset?resource=download.

**Ethics Approval:** Not applicable.

**Conflicts of Interest:** The authors declare no conflicts of interest.

## Nomenclature

| Symbol | Description |
|---|---|
| $x_i$ | Feature value for dimension $i$ |
| $x_i^{\text{norm}}$ | Normalized feature value |
| $N$ | Total number of transactions |
| $\lambda_i(t)$ | Instantaneous firing rate for feature $i$ |
| $\lambda_{\max}$ | Maximum firing rate |
| $v_j(t)$ | Membrane potential of neuron $j$ |
| $\tau_m$ | Membrane time constant |
| $v_{\text{rest}}$ | Resting membrane potential |
| $v_{\text{th}}$ | Firing threshold potential |
| $v_{\text{reset}}$ | Reset potential after spike |
| $R_m$ | Membrane resistance |
| $I_j(t)$ | Input current to neuron $j$ |
| $w_{ij}$ | Synaptic weight from neuron $i$ to $j$ |
| $t_i^f$ | Spike time from neuron $i$ |
| $\alpha(t)$ | Synaptic current kernel |
| $\tau_s$ | Synaptic time constant |
| $\phi_j(t)$ | Temporal receptive field of neuron $j$ |
| $s_j(t)$ | Spike indicator function |
| $\kappa(\tau)$ | Temporal weighting kernel |
| $C_{jk}(\Delta t)$ | Cross-correlation between neurons $j$ and $k$ |
| $A(t)$ | Anomaly score at time $t$ |
| $\theta_A$ | Anomaly detection threshold |
| $\Delta t$ | Time difference (pre-post spike) |
| $A_+, A_-$ | STDP learning rates (potentiation, depression) |
| $\tau_+, \tau_-$ | STDP time constants |
| $\eta$ | Learning rate |
| $w_{\max}$ | Maximum synaptic weight |
| $r_j$ | Firing rate of neuron $j$ |
| $r_{\text{target}}$ | Target firing rate (homeostasis) |
| $\tau_{\text{homeo}}$ | Homeostatic time constant |
| $p_{\text{conn}}$ | Connection probability |
| $V_{\text{consensus}}$ | Consensus voting score |
| $N_{\text{nodes}}$ | Number of blockchain nodes |
| $\theta_V$ | Consensus voting threshold |

| **Acronym** | **Full Form** |
|---|---|
| SNN | Spiking Neural Network |
| TSPR | Temporal Spike Pattern Recognition |

| | |
|---|---|
| STDP | Spike-Timing-Dependent Plasticity |
| LIF | Leaky Integrate-and-Fire |
| ROC | Receiver Operating Characteristic |
| AUC | Area Under Curve |
| MCC | Matthews Correlation Coefficient |
| FPR | False Positive Rate |
| FNR | False Negative Rate |

## References


1. Bäßler D, Kortus T, Gühring G. Unsupervised anomaly detection in multivariate time series with online evolving spiking neural networks. Mach Learn. 2022;111(4):3917–43. doi:10.1007/s10994-022-06129-4.
2. Bariah L, Shami A, Serhani MA. Anomaly detection in time series data using spiking neural networks. Adv Sci Lett. 2018;24(10):7728–31. doi:10.1166/asl.2018.12980.
3. Kumar A, Singh R. An efficient intrusion detection model based on convolutional spiking neural network. Sci Rep. 2024;14(1):6141. doi:10.1038/s41598-024-57691-x.
4. Roy K, Jaiswal A, Panda P. Towards spike-based machine intelligence with neuromorphic computing. Nature. 2019;575(7784):607–17. doi:10.1038/s41586-019-1677-2.
5. Kheradpisheh SR, Masquelier T. Unsupervised post-training learning in spiking neural networks. Sci Rep. 2025;15(1):10234. doi:10.1038/s41598-025-01749-x.
6. Mounnan O, Manad O, Boubchir L, El Mouatasim A, Daachi B. A review on deep anomaly detection in blockchain. Blockchain: Res Applicat. 2024;5(4):100227. doi:10.1016/j.bcra.2024.100227.
7. Hassan M, Rahman MS, Janicke H, Sarker IH. Detecting anomalies in blockchain transactions using machine learning classifiers and explainability analysis. Blockchain: Res Appl. 2024;5(3):100207.
8. Kousias N, Karamitsos I, Sharma PK. Anomaly detection in blockchain networks using unsupervised learning: a survey. Algorithms. 2024;17(5):201. doi:10.3390/a17050201.
9. Shevchuk R, Martsenyuk V, Adamyk B, Benson V, Melnyk A. Anomaly detection in blockchain: a systematic review of trends, challenges, and future directions. Appl Sci. 2025;15(15):8330. doi:10.3390/app15158330.
10. Ashfaq T, Khalid R, Yahaya AS, Aslam S, Azar AT, Alsafari S, et al. A machine learning and blockchain based efficient fraud detection mechanism. Sensors. 2022;22(19):7162. doi:10.3390/s22197162.
11. Esser SK, Merolla PA, Arthur JV, Cassidy AS, Appuswamy R, Modha DS, et al. Convolutional networks for fast, energy-efficient neuromorphic computing. Proc Nat Acad Sci. 2016;113(41):11441–6. doi:10.1073/pnas.1604850113.
12. Jumani F, Raza M. Machine learning for anomaly detection in blockchain: a critical analysis, empirical validation, and future outlook. Computers. 2025;14(7):247. doi:10.3390/computers14070247.
13. Bhende MS, Quraishi A, AlGhamdi A, Keshta I, Soni M, Singh BK, et al. A hybrid PKI and spiking neural network approach for enhancing security and energy efficiency in IoMT-based healthcare 5.0. SLAS Technol. 2025;32(2):100284. doi:10.1016/j.slast.2025.100284.
14. Mustafa YEA, Mustafa MMA, Eljack Babiker SM. Hybrid recurrent with spiking neural network model for enhanced anomaly prediction in IoT networks security. Front Artif Intell. 2025;8:1651516. doi:10.3389/frai.2025.1651516.
15. Vázquez IX, Sedano J, Afzal M, García-Vico ÁM. Vacuum spiker: a spiking neural network-based model for efficient anomaly detection in time series. arXiv:2510.06910. 2025.
16. Masquelier T, Thorpe SJ. Unsupervised learning of visual features through spike timing dependent plasticity. PLoS Comput Biol. 2007;3(2):e31. doi:10.1371/journal.pcbi.0030031.
17. Wu J, Chua Y, Zhang M, Li H, Tan KC. SSTDP: supervised spike timing dependent plasticity for efficient spiking neural network training. Front Neurosci. 2021;15:756876. doi:10.3389/fnins.2021.756876.
18. Srinivasan G, Roy K. Characterization of generalizability of spike timing dependent plasticity trained spiking neural networks. Front Neurosci. 2021;15:695357. doi:10.3389/fnins.2021.695357.
19. Shouval HZ, Wang SS, Wittenberg GM. Spike timing dependent plasticity: a consequence of more fundamental learning rules. Front Comput Neurosci. 2010;4:19. doi:10.3389/fncom.2010.00019.

20. Lee C, Panda P, Srinivasan G, Roy K. Deep spiking convolutional neural network trained with unsupervised spike-timing-dependent plasticity. IEEE Trans Cogn Dev Syst. 2019;11(3):384–94. doi:10.1109/TCDS.2018.2833071.
21. Taherkhani A, Belatreche A, Li Y, Cosma G, Maguire LP, McGinnity TM. A review of learning in biologically plausible spiking neural networks. Neural Netw. 2020;122(1):253–72. doi:10.1016/j.neunet.2019.09.036.
22. Li J, Wu B, Xia X, Liu X, Yi X, Zhang X. Unsupervised backdoor detection and mitigation for spiking neural networks. arXiv:2510.06629. 2025.
23. Liu Q, Zhao Z, Wen Y, Wang Y, Lombardi F. Low-power computing with neuromorphic engineering. Adv Intell Syst. 2021;3(2):2000150. doi:10.1002/aisy.202000150.
24. Enuganti PK, Sai BB, Prodromakis T, Roy O. Neuromorphic computing and applications: a topical review. WIREs Data Min Know Disc. 2025;15(3):e1570. doi:10.1002/widm.70014.
25. Chen Y, Zhang G, Liu F, Wu B, Deng Y, Gao D, et al. Revolutionizing neuromorphic computing with memristor-based artificial neurons. J Semicond. 2025;46(6):061301. doi:10.1088/1674-4926/24110006.
26. Liu S, Mohammadi N, Yi Y. Quantization-aware training of spiking neural networks for energy-efficient spectrum sensing on loihi chip. IEEE Trans Green Commun Netw. 2024;8(2):827–38. doi:10.1109/TGCN.2023.3337748.