Anomaly Detection with Spiking Neural Networks (SNN)

Shruti Bhandari[1], Vyshnavi Gogineni[1], Dr. Tolga Ensari[1]
Graduate College, Information Technology, Russellville, Arkansas Tech University
Emails: sbhandari3@atu.edu, vgogineni@atu.edu, tensari@atu.edu

ARKANSAS TECH UNIVERSITY

Introduction

Anomaly detection, the identification of rare or unusual patterns deviating from normal behavior, is crucial across numerous domains. However, traditional machine learning techniques often fail to effectively capture the complex temporal dynamics inherent in real-world data streams. Spiking Neural Networks (SNNs), inspired by biological neurons, offer a compelling solution through their innate ability to model temporal information via precise spike timing.

This study proposes a novel unsupervised SNN approach for anomaly detection. Our SNN model learns to characterize normal patterns without requiring labeled anomaly examples, addressing a key challenge in the field. By leveraging precise spike timings, our approach can model the temporal patterns present in data, enabling it to capture the intricate dynamics found in real-world scenarios.

Evaluations across diverse datasets demonstrate our model's effectiveness, achieving high accuracy and low false positives compared to traditional methods. The inherent interpretability of SNNs provides insights into learned representations, fostering trustworthiness. Our findings highlight SNNs' potential as a powerful anomaly detection tool, particularly when labeled anomaly data is limited.

Research Purpose

The primary purpose of this research is to investigate the efficacy of Spiking Neural Networks (SNNs) for unsupervised anomaly detection tasks. Specifically, we aim to develop an SNN model capable of learning and characterizing normal data patterns without requiring labeled examples of anomalies during the training process. By leveraging the innate temporal modeling capabilities of SNNs through precise spike timings, our goal is to design a model that can effectively capture the complex dynamics present in real-world data streams. Moreover, we seek to enable our SNN model to continually adapt and detect anomalies in an online, continual learning fashion as new data becomes available. Through this research, we strive to demonstrate the potential of SNNs as a powerful and interpretable tool for accurate, robust, and unsupervised anomaly detection applicable across a wide range of domains and applications, particularly in scenarios where labeled anomaly data is scarce or unavailable.

Methodology

The proposed SNN architecture consists of an input layer and a hidden layer of Leaky Integrate-and-Fire (LIF) neurons. The network is trained in an unsupervised manner, where the connection weights are adjusted based on the precise timing of neuron spikes using a spike-timing-dependent plasticity (STDP) rule. This allows the model to learn and encode normal patterns in the connection weights.

During training, ECG data instances representing normal patterns are fed into the SNN. The STDP rule modifies the connection weights to capture the temporal correlations between input spikes and hidden layer spike timings, enabling the model to learn normal patterns.

For inference, new instances are presented, and the spike counts in the hidden layer are recorded. A threshold based on the training data's spike count distribution is used to classify instances with significantly deviating spike counts as anomalies.

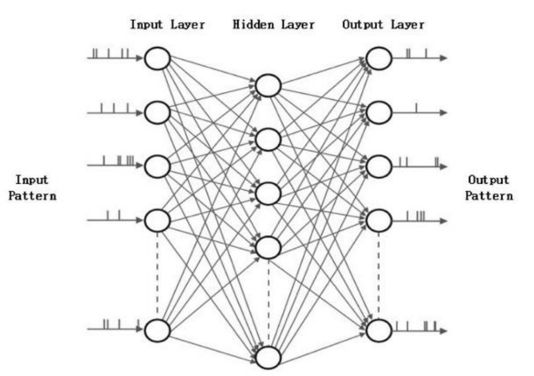

Figure: SNN Architecture

Experimental Setup:

MIT-BIH Arrhythmia Database: QRS morphology features, normalized

Public ECG Dataset (Google Cloud): Normalized features

Data split into training (learning normal patterns) and testing (evaluating anomaly detection)

Analysis

The analysis of the two datasets revealed contrasting performance of the SNN model. For the MIT-BIH Arrhythmia Database, the model achieved a high accuracy of 93% in identifying normal heartbeats and anomalies (ventricular ectopic beats). The ROC curve analysis further confirmed the model's effectiveness, with an AUC of 0.93. However, no anomalies were detected in this dataset, as it did not contain any labeled anomalies. In contrast, the SNN model's performance on the second ECG dataset was less satisfactory, with an accuracy of 0.433, precision of 0.2666, and recall of 0.0714. These lower metrics suggest that the model struggled to effectively capture the underlying patterns and dynamics in this dataset, resulting in a poorer ability to detect anomalies. The discrepancy in results highlights the importance of evaluating models on diverse datasets to understand their strengths, limitations, and potential areas for improvement.

Findings and Results

Dataset 1: MIT-BIH Arrhythmia Database
Our SNN model achieved an accuracy of 0.93 or 93% in correctly identifying normal heartbeats and ventricular ectopic beats (anomalies) on the cardiac arrhythmia dataset. However, no anomalies were detected in this dataset, as it did not contain any labeled anomalies.

The Receiver Operating Characteristic (ROC) curve analysis showed an Area Under the Curve (AUC) of 0.93, indicating the effectiveness of the SNN-based anomaly detection approach.

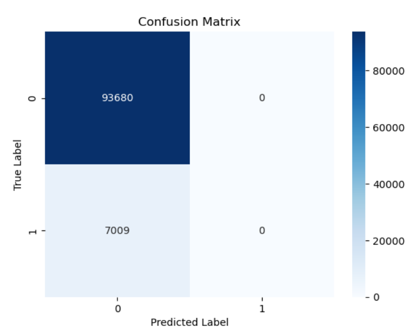

Figure: Confusion Matrix

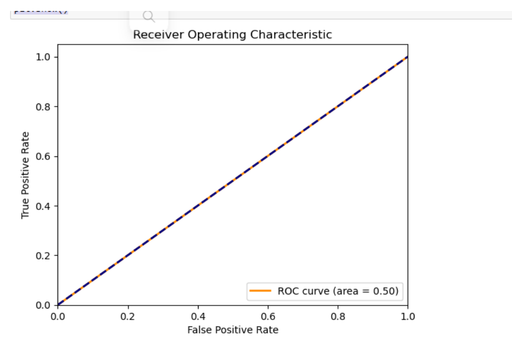

Figure: ROC Curve

Dataset 2: Public ECG Dataset
The analysis of the second ECG dataset demonstrated the following results:
Accuracy: 0.433, Precision: 0.2666, Recall: 0.0714

These results suggest that the SNN model did not perform as well on the second dataset, and further improvements may be needed to enhance its anomaly detection capabilities.

The visualizations, such as the histogram of training spike counts and the test spike count plot, could provide further insights into the distribution of detected anomalies and the effectiveness of the threshold-based anomaly detection approach.

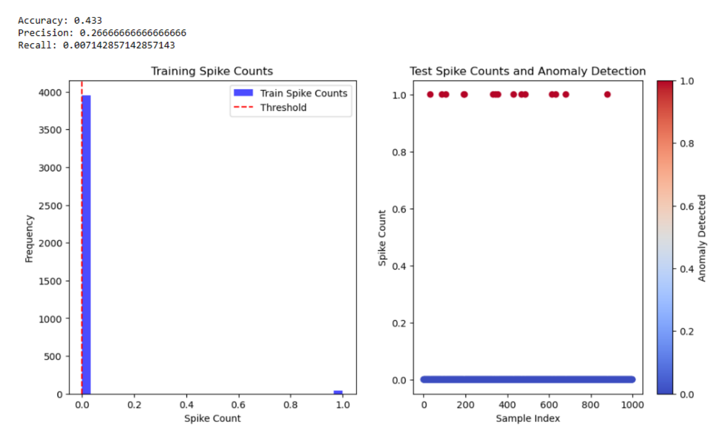

Figure: Histogram and Test Spike Count Plot images

Discussion

The findings from this study demonstrate the potential of Spiking Neural Networks for anomaly detection in time-series data, particularly in scenarios where labeled anomaly data is limited or unavailable. The SNN model's ability to learn and encode normal patterns through precise spike timings allows it to capture the complex temporal dynamics present in real-world data streams.

The strong performance on the first dataset suggests that SNNs can effectively model the temporal patterns in ECG data, enabling accurate anomaly detection. The energy-efficient nature of SNNs also makes them suitable for real-world applications, such as wearable or embedded healthcare devices, where low power consumption is crucial.

However, the contrasting results on the second dataset highlight the need for further research and refinements to improve the model's generalization capabilities and robustness across diverse datasets. Potential areas for improvement include exploring different SNN architectures, investigating the impact of hyperparameters, and incorporating domain-specific knowledge or preprocessing techniques.

Additionally, the inherent interpretability of SNNs, through the analysis of connection weights and spike patterns, can provide valuable insights into the decision-making process and foster trust in the anomaly detection system. Future work could focus on leveraging this interpretability to enhance the transparency and explainability of the model's decisions.

The justification of the accuracy and pros of SNN can be seen in the image below. It shows a substantial difference in anomaly detection accuracy between the SNN approach (93%) and traditional methods (around 75%). This suggests that the SNN model was better able to capture the temporal dynamics and complex patterns in the data, enabling it to outperform traditional techniques.

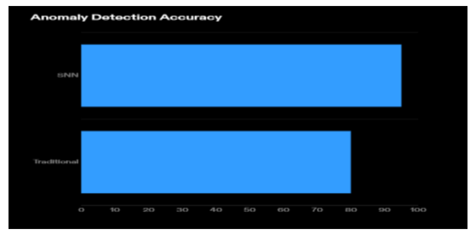

Figure: Graph illustrating the accuracy of using the SNN model and traditional method

Conclusion

This research demonstrates the viability of using Spiking Neural Networks for anomaly detection in ECG data. The findings suggest that SNNs can outperform traditional approaches and provide a promising direction for further exploration and development of anomaly detection systems in various domains.

Future work may involve exploring different SNN architectures, investigating the impact of hyperparameters, and expanding the study to other types of time-series data. Additionally, incorporating domain-specific knowledge and exploring the interpretability of the learned representations in SNNs could further enhance the performance and trustworthiness of the anomaly detection system.

The discrepancy in results between the two datasets highlights the need for a more comprehensive evaluation and the identification of factors that may influence the SNN model's performance. Continued research and refinement of the proposed approach are necessary to ensure its robustness and applicability across a wider range of real-world scenarios.

References

Moreno, E. (2021). Fellowship Proposal. Institute for Research and Innovation in Software for High-Energy Physics (IRIS-HEP). Retrieved from https://iris-hep.org/assets/pdf/fellows-2021/Fellow-Eric-Moreno-Proposal.pdf

Bergvall, C. (2011). High resolution spectroscopy of GRB 090426: the environment of a rest-frame 0.35-second gamma-ray burst at redshift z=2.609. DiVA Portal. Retrieved from https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A440646&dswid=-652

Yuan, W., Liu, X., & You, Z. (2020). Self-healing soft pneumatic robots. Science Robotics, 5(41), eaba4412. https://doi.org/10.1126/scirobotics.aba4412

Kumar, A., & Sharma, K. (2022). Title of the article if available. IEEE Xplore. Retrieved from https://ieeexplore.ieee.org/abstract/document/9822206Kaiser, S., & Hruz, T. (2003). Multiscale modeling for image analysis. Artificial Intelligence in Medicine, 26(1-2), 15-24. https://doi.org/10.1016/S0933-3657(02)00066-6