Adversarial Robustness and Inference Latency Trade-offs in BE-SNNs vs. Standard SNNs

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: What is the trade-off between adversarial robustness and inference latency in BE-SNNs versus standard SNNs when tested on adversarial samples from tabular datasets, measured in throughput and accuracy. 7 claims were extracted from source literature; 4 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Enhancing Adversarial Robustness in SNNs with Sparse Gradients. Research question: What is the trade-off between adversarial robustness and inference latency in BE-SNNs versus standard SNNs when tested on adversarial samples from tabular datasets, measured in throughput and accuracy?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.7/10.

3 Results

11 papers retrieved. 7 claims extracted; 4 independently verified. Quality review score: 6.7/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
The random vulnerability of f at point x to an p attack of size ϵ is defined as the expected value of (fy(x + ϵ $\cdot$ δ) -f	×	0.01
The adversarial vulnerability of f at point x to an p attack of size ϵ is defined as the supremum of (fy(x+$\epsilon$$\cdot$$\delta$)-fy(x))2	×	0.01
Adversarial examples generated under ∞ attacks tend to be more destructive compared to those generated under 0 and 2	×	0.04
SNNs exhibit greater resilience to random perturbations compared to adversarial perturbations, even at larger scales.	✓	0.30
The performance gap between SNNs under adversarial and random perturbations is upper bounded by the gradient sparsity of	✓	0.38
The proposed gradient sparsity regularization strategy improves the robustness of SNNs.	✓	0.20
The effectiveness of the proposed approach is validated through extensive experiments on both image-based and event-base	✓	0.18

References

- http://arxiv.org/abs/2311.10802v1
- http://arxiv.org/abs/2504.20900v1
- http://arxiv.org/abs/2405.20355v1