Stochastic Sampling in Bitwise Neural Networks Enhances Adversarial Robustness on MNIST

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 10 peer-reviewed papers addressing the following research question: Does stochastic sampling in bitwise neural networks improve robustness against adversarial attacks on MNIST compared to deterministic weight-based BNNs, as measured by adversarial accuracy under FGSM. 15 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.4/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Robust Image Classification: Defensive Strategies against FGSM and PGD Adversarial Attacks. Research question: Does stochastic sampling in bitwise neural networks improve robustness against adversarial attacks on MNIST compared to deterministic weight-based BNNs, as measured by adversarial accuracy under FGSM and PGD attacks?.

2 Methodology

Systematic literature search across multiple databases yielded 10 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.4/10.

3 Results

10 papers retrieved. 15 claims extracted; 1 independently verified. Quality review score: 4.4/10.

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 concepts of adversarial examples and the FGSM attack were introduced in reference [1].	×	0.09
Adversarial training can be computationally expensive and may not generalize well to unseen types of attacks.	×	0.05
PGD was proposed in reference [2] as a robust method for generating adversarial examples and for use in adversarial trai	×	0.09
Adversarial training with PGD significantly enhances the robustness of deep learning models.	✓	0.15
Reference [6] introduces sophisticated attacks that successfully bypass ten state-of-the-art detection methods.	×	0.06
Reference [6] does not propose any improved detection mechanisms for the attacks it presents.	×	0.04
Reference [7] proposes a novel adversarial attack targeting image captioning models using attention-based optimization t	×	0.06
On the MNIST dataset, the model accuracy drops from 0.9927 at noise level 0.00 to 0.0122 at noise level 0.30.	×	0.02
On the MNIST Fashion dataset, the model accuracy remains relatively stable between 0.2943 and 0.2956 for noise levels ra	×	0.01
In the second experiment table, MNIST accuracy at noise level 0.00 is 0.9909, dropping to 0.0528 at noise level 1.00.	×	0.01
For the defense mechanism tested in Table (p7), the time required to defend an attack on MNIST is 0.003 seconds at noise	×	0.02
At noise level 0.10, the test accuracy on MNIST with the first defense method is 0.9372, achieved in 0.0003 seconds.	×	0.06
At noise level 1.00, the test accuracy on MNIST Fashion with the first defense method is 0.4605.	×	0.03
In the second defense experiment table, the time for defending an attack is consistently 0.0012 seconds across all teste	×	0.02
At noise level 0.05, the test accuracy on MNIST in the second defense experiment is 0.9342.	×	0.02

References

- http://arxiv.org/abs/2408.13274v1
- http://arxiv.org/abs/2307.02055v1
- http://arxiv.org/abs/1611.06539v1