SOVEREIGN: How does the performance of NOVA's open-world recognition capability compare to existing OOD detection methods

SOVEREIGN Research Kernel
Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Mahalanobis distance (MD) is a simple and popular post-processing method for detecting out-of-distribution (OOD) inputs in neural networks. We analyze its failure modes for near-OOD detection and propose a simple fix called relative Mahalanobis distance (RMD) which improves performance and is more robust to hyperparameter choice. On a wide selection of challenging vision, language, and biology OOD benchmarks (CIFAR-100 vs CIFAR-10, CLINC OOD intent detection, Genomics OOD), we show that RMD meaningfully improves upon MD performance (by up to 15% AUROC on genomics OOD).

1 Introduction

Analysis of: A Simple Fix to Mahalanobis Distance for Improving Near-OOD Detection. Research goal: How does the performance of NOVA's open-world recognition capability compare to existing OOD detection methods like ODIN and Mahalanobis distance when evaluated on standard benchmarks such as CIFAR-100 and ImageNet-30?.

2 Methodology

Multi-query arXiv search (4 parallel queries, Relevance-sorted). TF-IDF cosine semantic verification (bigrams, threshold=0.15). NIM nv-embedqa-e5-v5 (dim=1024) for semantic indexing. Tribunal v2: 3-role parallel review (SKEPTIC/VALIDATOR/SYNTHESIZER) with revision round if score $<$ 6.5.

3 Results

5 papers retrieved. 4 claims extracted, 4 verified. Tribunal: 7.8/10 → APPROVE (revision_round=0). Policy: AUTO_APPROVE.

4 Uncertainties

NIM free tier latency varies. TF-IDF verification is a weak signal. arXiv Relevance ranking is query-dependent. Tribunal consensus is LLM-based and prompt-sensitive.

5 Extracted Claims

Claim	Verified	Confidence
Mahalanobis distance is a simple and popular post-processing method for detecting out-of-distribution inputs in neural n	✓	0.36
Relative Mahalanobis distance improves performance and is more robust to hyperparameter choice compared to Mahalanobis d	✓	0.28
Relative Mahalanobis distance meaningfully improves upon Mahalanobis distance performance by up to 15% AUROC on genomics	✓	0.35
The evaluation was conducted on CIFAR-100 vs CIFAR-10, CLINC OOD intent detection, and Genomics OOD benchmarks.	✓	0.39

References

- http://arxiv.org/abs/2106.09022v1
- http://arxiv.org/abs/astro-ph/0509589v1
- http://arxiv.org/abs/1604.02275v1