

SOVEREIGN: How does the MambaFormer hybrid MoE architecture’s efficiency (FLOPs per token and throughput) scale with mode

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Multilayer neural networks trained with the back-propagation algorithm constitute the best example of a successful gradient based learning technique. Given an appropriate network architecture, gradient-based learning algorithms can be used to synthesize a complex decision surface that can classify high-dimensional patterns, such as handwritten characters, with minimal preprocessing. This paper reviews various methods applied to handwritten character recognition and compares them on a standard handwritten digit recognition task. Convolutional neural networks, which are specifically designed to

1 Introduction

Analysis of: Gradient-based learning applied to document recognition. Research goal: How does the MambaFormer hybrid MoE architecture’s efficiency (FLOPs per token and throughput) scale with model width and number of experts on code generation tasks using HumanEval and MBPP, relative to standard Transformer MoE models?.

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. 1 claims extracted, 1 verified. Tribunal: 7.2/10 → 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
Multilayer neural networks trained with the back-propagation algorithm constitute the best example of a successful gradi	✓	0.34

References

- <https://doi.org/10.1109/5.726791>
- <https://doi.org/10.1007/s11704-026-60308-3>
- <https://doi.org/10.18653/v1/2020.coling-main>