

# SOVEREIGN: What is the impact of cross-modal expert routing consistency on SEED-Bench zero-shot accuracy when varying the

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 29, 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: What is the impact of cross-modal expert routing consistency on SEED-Bench zero-shot accuracy when varying the number of visual tokens in MoE VLMs versus dense baselines?.

## 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

11 papers retrieved. 0 claims extracted, 0 verified. Tribunal: 6.7/10 → APPROVE (revision\_round=1). Policy: ESCALATE\_TO\_OWNER.

## 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.

## References

- <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- <https://doi.org/10.1109/5.726791>
- <https://doi.org/10.1007/s11704-026-60308-3>