

SOVEREIGN: How does AnyExperts' on-demand routing strategy compare to fixed routing baselines in terms of inference latency

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Sparse Mixture-of-Experts (MoE) architectures enable efficient scaling of large language models through conditional computation, yet the routing mechanisms responsible for expert selection remain poorly understood. In this work, we introduce routing signatures, a vector representation summarizing expert activation patterns across layers for a given prompt, and use them to study whether MoE routing exhibits task-conditioned structure. Using OLMoE-1B-7B-0125-Instruct as an empirical testbed, we show that prompts from the same task category induce highly similar routing signatures, while prompts

1 Introduction

Analysis of: Task-Conditioned Routing Signatures in Sparse Mixture-of-Experts Transformers. Research goal: How does AnyExperts' on-demand routing strategy compare to fixed routing baselines in terms of inference latency and throughput across varying batch sizes on standard multimodal benchmarks like VQAv2 and Visual7W?

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. 6 claims extracted, 6 verified. Tribunal: 8.5/10 → 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
Prompts from the same task category induce highly similar routing signatures, while prompts from different categories ex	✓	0.33
Within-category routing similarity (0.8435 +/- 0.0879) significantly exceeds across-category similarity (0.6225 +/- 0.16	✓	0.29
A logistic regression classifier trained solely on routing signatures achieves 92.5% +/- 6.1% cross-validated accuracy o	✓	0.31
The observed separation in routing signatures is not explained by sparsity or load-balancing constraints alone.	✓	0.20
Task structure becomes increasingly apparent in deeper layers of the MoE transformer.	✓	0.20
Routing in sparse transformers is not merely a balancing mechanism, but a measurable task-sensitive component of conditi	✓	0.30

References

- <http://arxiv.org/abs/1007.0409v1>
- <http://arxiv.org/abs/2603.11114v1>
- <http://arxiv.org/abs/1403.3007v3>