SOVEREIGN: What is the inference efficiency tradeoff between SMoES and hard-routing MoE approaches when evaluated on lang

SOVEREIGN Research Kernel
Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Abstract Large language models (LLMs) have demonstrated impressive capabilities, but the bar for clinical applications is high. Attempts to assess the clinical knowledge of models typically rely on automated evaluations based on limited benchmarks. Here, to address these limitations, we present MultiMedQA, a benchmark combining six existing medical question answering datasets spanning professional medicine, research and consumer queries and a new dataset of medical questions searched online, HealthSearchQA. We propose a human evaluation framework for model answers along multiple axes including

1 Introduction

Analysis of: Large language models encode clinical knowledge. Research goal: What is the inference efficiency tradeoff between SMoES and hard-routing MoE approaches when evaluated on language model reasoning tasks across varying input modalities?.

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

8 papers retrieved. 7 claims extracted, 7 verified. Tribunal: 7.5/10 $\rightarrow$ 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
Large language models (LLMs) have demonstrated impressive capabilities, but the bar for clinical applications is high.	✓	0.28
MultiMedQA is a benchmark combining six existing medical question answering datasets spanning professional medicine, res	✓	0.37
Flan-PaLM achieves state-of-the-art accuracy on every MultiMedQA multiple-choice dataset (MedQA, MedMCQA, PubMedQA and M	✓	0.34
Flan-PaLM achieves 67.6% accuracy on MedQA (US Medical Licensing Exam-style questions), surpassing the prior state of th	✓	0.33
Human evaluation reveals key gaps in Flan-PaLM's performance.	✓	0.20
Instruction prompt tuning is a parameter-efficient approach for aligning LLMs to new domains using a few exemplars.	✓	0.28
The resulting model, Med-PaLM, performs encouragingly, but remains inferior to clinicians.	✓	0.24

References

- https://doi.org/10.1038/s41586-023-06291-2
- https://doi.org/10.1109/tmi.2014.2377694
- https://doi.org/10.1007/s11704-026-60308-3