

SOVEREIGN: What is the robustness of SMOES-based MoE-VLMs to distribution shift and adversarial inputs compared to dense

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Language models demonstrate both quantitative improvement and new qualitative capabilities with increasing scale. Despite their potentially transformative impact, these new capabilities are as yet poorly characterized. In order to inform future research, prepare for disruptive new model capabilities, and ameliorate socially harmful effects, it is vital that we understand the present and near-future capabilities and limitations of language models. To address this challenge, we introduce the Beyond the Imitation Game benchmark (BIG-bench). BIG-bench currently consists of 204 tasks, contributed b

1 Introduction

Analysis of: Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. Research goal: What is the robustness of SMOES-based MoE-VLMs to distribution shift and adversarial inputs compared to dense VLMs of equivalent inference FLOPs on multimodal benchmarks?.

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

10 papers retrieved. 4 claims extracted, 4 verified. Tribunal: 8.2/10 → 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
BIG-bench consists of 204 tasks contributed by 450 authors across 132 institutions.	✓	0.26
BIG-bench focuses on tasks believed to be beyond the capabilities of current language models.	✓	0.30
Model performance and calibration both improve with scale, but are poor in absolute terms and when compared with rater p	✓	0.25
Performance is remarkably similar across model classes, though with benefits from sparsity.	✓	0.19

References

- <https://doi.org/10.48550/arxiv.2206.04615>
- <https://doi.org/10.1609/aaai.v32i1.11671>
- <https://doi.org/10.1145/2523813>