SOVEREIGN: What is the impact of expert capacity imbalance on AnyExperts' performance degradation when evaluated on domai

SOVEREIGN Research Kernel
Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Vision-language foundation models achieve promising performance in natural image classification, yet their direct application to medical imaging is limited by severe domain shifts, resolution mismatches, and the multi-label nature of clinical diagnosis. Training dedicated medical foundation models from scratch, however, is costly and data-intensive. Here, we propose MedBridge, a lightweight adaptation framework that opens a new direction in domain-gap mitigation by jointly combining domain alignment, resolution preservation, and multi-label reasoning via complementary VLM experts for medical i

1 Introduction

Analysis of: Adapting Foundation Vision-Language Models to Medical Diagnosis via Query-Driven Expert Bridging. Research goal: What is the impact of expert capacity imbalance on AnyExperts' performance degradation when evaluated on domain-shifted datasets such as migrating from COCO-based training to Conceptual Captions testing?.

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

1 papers retrieved. 10 claims extracted, 10 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
Vision-language foundation models achieve promising performance in natural image classification.	✓	0.29
Direct application of vision-language foundation models to medical imaging is limited by severe domain shifts, resolutio	✓	0.39
Training dedicated medical foundation models from scratch is costly and data-intensive.	✓	0.26
MedBridge is a lightweight adaptation framework that jointly combines domain alignment, resolution preservation, and mul	✓	0.36
MedBridge transforms pretrained VLMs into multi-view query encoders that inject a compact set of learnable query tokens	✓	0.34
MedBridge enables non-destructive domain alignment while preserving fine-grained pathological cues via multi-view high-r	✓	0.31
Query tokens in MedBridge act as routing signals for a mixture-of-experts, dynamically integrating heterogeneous foundat	✓	0.38
MedBridge was evaluated on five chest radiograph benchmarks in three key adaptation tasks.	✓	0.19
MedBridge demonstrates superior performance in both cross-domain generalization (out-of-distribution transfer) and in-do	✓	0.31
MedBridge yields a significant 6-15% AUC improvement over state-of-the-art adaptation methods for multi-label thoracic d	✓	0.27

References

- https://www.semanticscholar.org/paper/dad4071465b4b2a18a1589c93abc3328de96231f