

SOVEREIGN: Adapting Foundation Vision-Language Models to Medical Diagnosis via Query-Driven

SOVEREIGN Research Kernel
Autonomous draft — Owner review required before publication

May 27, 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: How does the performance of dynamic expert specialization in MoE-VLMs degrade under distribution shift compared to fixed routing strategies on out-of-distribution multimodal reasoning 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

4 papers retrieved. 5 claims extracted, 4 verified. Tribunal: 7.2/10 → 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
MedBridge is a lightweight adaptation framework for medical image diagnosis.	✓	0.21
MedBridge transforms pretrained VLMs into multi-view query encoders that inject learnable query tokens into intermediate	✓	0.29
MedBridge uses query tokens as routing signals for a mixture-of-experts to integrate heterogeneous foundation models.	✓	0.20
MedBridge was evaluated on five chest radiograph benchmarks.	×	0.11
MedBridge yields a 6-15% AUC improvement over state-of-the-art adaptation methods for multi-label thoracic disease diagn	✓	0.26

References

- <https://www.semanticscholar.org/paper/b79af713b2b7ca628c569c30f0a814a28f10086b>
- <https://www.semanticscholar.org/paper/dad4071465b4b2a18a1589c93abc3328de96231f>
- <https://www.semanticscholar.org/paper/e27931b44157478a13c15d501d3d5bb0830f38df>