

Adversarial Pretraining Strategies for Robust Multimodal Language Models Under Distribution Shifts

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: Can adversarial pretraining strategies for tabular foundation models be adapted to improve the robustness of multimodal language models against distribution shifts in vision-language benchmarks. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Toward a Holistic Evaluation of Robustness in CLIP Models. Research question: Can adversarial pretraining strategies for tabular foundation models be adapted to improve the robustness of multimodal language models against distribution shifts in vision-language benchmarks?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.2/10.

3 Results

15 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 5.2/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

References

- <http://arxiv.org/abs/2410.01534v2>
- <http://arxiv.org/abs/2512.03307v1>
- <http://arxiv.org/abs/2403.09513v1>