

Adversarial Robustness of Qwen14B, LLaVA-1.5, and InstructBLIP on GSM8K-V under Visual Noise Perturbations

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: How does the adversarial robustness of Qwen14B on GSM8K-V compare to LLaVA-1.5 and InstructBLIP when fine-tuned with different alignment techniques under varying levels of visual noise perturbation. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: On the Adversarial Robustness of Vision Transformers. Research question: How does the adversarial robustness of Qwen14B on GSM8K-V compare to LLaVA-1.5 and InstructBLIP when fine-tuned with different alignment techniques under varying levels of visual noise perturbation?.

2 Methodology

Systematic literature search across multiple databases yielded 16 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.7/10.

3 Results

16 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.7/10.

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/2503.14504v2>
- <http://arxiv.org/abs/2507.22398v3>
- <http://arxiv.org/abs/2103.15670v3>