Qwen14B Performance on GSM8K-V Under Adversarial Visual Noise: A Multimodal Model Comparison

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does the performance of Qwen14B on GSM8K-V compare to other multimodal models like LLaVA-1.5 and InstructBLIP when trained with different levels of adversarial visual noise, and what is the. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: GSM8K-V: Can Vision Language Models Solve Grade School Math Word Problems in Visual Contexts. Research question: How does the performance of Qwen14B on GSM8K-V compare to other multimodal models like LLaVA-1.5 and InstructBLIP when trained with different levels of adversarial visual noise, and what is the degradation rate in accuracy?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

14 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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/2509.25160v1
- http://arxiv.org/abs/2407.04973v1
- http://arxiv.org/abs/2308.07921v1