

# Adversarial Visual Noise Impact on Qwen14B GSM8K-V Performance vs. LLaVA-1.5 and InstructBLIP

Assignee Research

June 7, 2026

## Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: How does the performance of Qwen14B on GSM8K-V degrade under adversarial visual noise compared to LLaVA-1.5 and InstructBLIP. 18 claims were extracted from source literature; 3 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Enhancing Visual Question Answering through Ranking-Based Hybrid Training and Multimodal Fusion. Research question: How does the performance of Qwen14B on GSM8K-V degrade under adversarial visual noise compared to LLaVA-1.5 and InstructBLIP?.

## 2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.8/10.

## 3 Results

13 papers retrieved. 18 claims extracted; 3 independently verified. Quality review score: 4.8/10.

## 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.



## 5 Extracted Claims

Claim	Verified	Confidence
The RankVQA model was evaluated using two major visual question answering (VQA) datasets: VQA v2.0 and COCO-QA.	✓	0.20
VQA v2.0 contains over 200,000 images and 600,000 questions.	×	0.06
COCO-QA comprises 123,287 images and over 117,000 questions.	×	0.05
The experimental environment for training and evaluating the RankVQA model included an NVIDIA Tesla V100 (32GB) GPU.	×	0.03
The experimental environment included an Intel Xeon E5-2698 v4 CPU.	×	0.00
The experimental environment included 256GB DDR4 memory.	×	0.01
The experimental environment included 2TB SSD storage.	×	0.01
The experimental environment used Ubuntu 20.04 LTS as the operating system.	×	0.01
The experimental environment used PyTorch 1.10.0 as the deep learning framework.	×	0.04
The experimental environment used CUDA Version 11.2.	×	0.02
The experimental environment used cuDNN Version 8.1.	×	0.02
The experimental environment used Python Version 3.8.10.	×	0.03
During data preprocessing, all images were resized to a uniform size of 224x224 pixels.	×	0.02
During data preprocessing, images were normalized by scaling the pixel values to the range of 0 to 1.	×	0.03
The RankVQA model employs the Faster R-CNN model to extract visual features from images.	×	0.12
The RankVQA model utilizes a pre-trained BERT model to extract text features.	×	0.14
The RankVQA model uses a multi-head self-attention mechanism for multimodal fusion.	✓	0.15
The RankVQA model includes a ranking learning module to improve answering accuracy by optimizing the relative ranking of	✓	0.16

## References

- <http://arxiv.org/abs/2305.14882v2>
- <http://arxiv.org/abs/2509.25160v1>
- <http://arxiv.org/abs/2408.07303v2>