Multi-Modal Embedding Integration in Blended RAG for Multimodal QA Accuracy

Assignee Research

May 31, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: What is the impact of incorporating multi-modal embeddings (e.g., CLIP for images and text) in the hybrid retrieval process of Blended RAG on accuracy for multimodal QA benchmarks like MMBench and. Retrieval-Augmented Generation (RAG) has been introduced to mitigate hallucinations in Multimodal Large Language Models (MLLMs) by incorporating external knowledge into the generation process, and it has become a widely adopted approach for knowledge-intensive Visual Question. 10 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: QA-Dragon: Query-Aware Dynamic RAG System for Knowledge-Intensive Visual Question Answering. Research question: What is the impact of incorporating multi-modal embeddings (e.g., CLIP for images and text) in the hybrid retrieval process of Blended RAG on accuracy for multimodal QA benchmarks like MMBench and LLaVA-Bench?.

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 4.0/10.

3 Results

15 papers retrieved. 10 claims extracted; 0 independently verified. Quality review score: 4.0/10.

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
QA-Dragon achieves an accuracy of 21.31% in single-source tasks and 23.22% in multi-source tasks.	×	0.12
Removing the domain router leads to a drop in accuracy to 19.04% in single-source tasks and 21.25% in multi-source tasks	×	0.08
Disabling the tool router results in a performance degradation to 18.32% in single-source tasks.	×	0.04
Eliminating query splitting significantly degrades performance.	×	0.04
Removing two-stage reranking decreases accuracy to 20.90% in single-source tasks and 22.14% in multi-source tasks, and i	×	0.05
The framework achieves the best overall balance by dynamically coordinating reasoning and retrieval.	×	0.06
QA-Dragon decomposes the problem into three branches with multiple processes: a Pre-Answer Module, a Search Router, and	×	0.10
The Pre-Answer Module performs domain classification and generates an initial reasoning trace and answer using a domain-	×	0.08
The Search Router inspects the reasoning trace to determine whether additional external evidence is required, and if so,	×	0.09
The Tool Router decides whether to invoke an Image Search Agent or a Text Search Agent.	×	0.07

References

- http://arxiv.org/abs/2510.20193v1
- http://arxiv.org/abs/2508.05197v2
- http://arxiv.org/abs/2404.07220v2