

Retrieval-Augmented Generation vs Zero-Shot Reranking Latency on TriviaQA at Scale

Assignee Research

June 2, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: How does the inference latency of retrieval-augmented generation compare to zero-shot re-ranking methods on the TriviaQA benchmark when scaling to larger context windows. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: FlashSpeech: Efficient Zero-Shot Speech Synthesis. Research question: How does the inference latency of retrieval-augmented generation compare to zero-shot re-ranking methods on the TriviaQA benchmark when scaling to larger context windows?.

2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.7/10.

3 Results

12 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 4.7/10.

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/2404.14700v4>
- <http://arxiv.org/abs/2502.00306v2>
- <http://arxiv.org/abs/2402.12317v2>