

How does the retrieval efficiency of Llama-3-8B-128K vary across context lengths 32K, 64K, and 128K on the MuSiQue benchmark when using the Tree of Reviews framework compared to chain-based retrieval

Assignee Research

May 29, 2026

Abstract

Recent advancements in Large language models (LLMs) have demonstrated remarkable capabilities across diverse domains. While they exhibit strong zero-shot performance on various tasks, LLMs' effectiveness in music-related applications remains limited due to the relatively small proportion of music-specific knowledge in their training data. To address this limitation, we propose Must-RAG, a comprehensive framework based on Retrieval Augmented Generation (RAG) to adapt general-purpose LLMs for text-only music question answering (MQA) tasks. RAG is a technique that provides external knowledge to

1 Introduction

This paper examines: MUST-RAG: MUSical Text Question Answering with Retrieval Augmented Generation. Research question: How does the retrieval efficiency of Llama-3-8B-128K vary across context lengths 32K, 64K, and 128K on the MuSiQue benchmark when using the Tree of Reviews framework compared to chain-based retrieval?.

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. 15 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.

5 Extracted Claims

Claim	Verified	Confidence
The evaluation used two datasets: ArtistMus (in-domain) and TrustMus (out-of-domain).	×	0.11
Performance on factual and contextual questions was separately measured on the ArtistMus dataset.	×	0.03
TrustMus evaluation was conducted across four categories: People (Ppl), Instrument & Technology (IT), Genre, Forms, and	×	0.02
All evaluations use a multiple-choice QA format.	×	0.02
Zero-shot baselines evaluated include GPT-4o, Llama 3.1 8B Instruct, MuLLaMA, and ChatMusician.	×	0.04
MuLLaMA is designed to handle audio-based question answering.	×	0.07
ChatMusician specializes in music understanding and generation with ABC notation.	×	0.05
Llama 3.1 8B Instruct was fine-tuned on 8K multiple-choice QA pairs generated from MusWikiDB.	×	0.03
RAG inference was implemented using Llama 3.1 8B Instruct and MusWikiDB as the retrieval database.	×	0.07
RAG fine-tuning was performed using a dataset in the form of (context, question, answer).	×	0.09
Models were trained for one epoch using LoRA with 8-bit quantization and specific hyperparameters.	×	0.03
For the ArtistMus dataset, half of the artists were included in the training data (Seen), while the other half were excl	×	0.04
MusWikiDB was developed by collecting music-related content from Wikipedia across seven categories: artists, genres, ins	×	0.06
MusWikiDB contains 31K pages, 629.2K passages, 65.5M total tokens, and a vocabulary size of 786K.	×	0.01
Wikipedia Corpus contains 3.2M pages, 21M passages, 2.1B total tokens, and a vocabulary size of 21.5M.	×	0.01

References

- <http://arxiv.org/abs/2507.23334v2>
- <http://arxiv.org/abs/2503.01763v2>
- <http://arxiv.org/abs/2404.14464v1>