

ArtistMus: A Globally Diverse, Artist-Centric Benchmark for Retrieval-Augmented Music Question Answering

Daeyong Kwon, Seungheon Doh, Juhan Nam*

Graduate School of Culture Technology, KAIST

Daejeon, South Korea

{ejmj63, seungheondoh, juhan.nam}@kaist.ac.kr

Abstract

Recent advances in Large Language Models (LLMs) have transformed open-domain question answering, yet their effectiveness in music-related reasoning remains limited due to sparse music knowledge in pretraining data. While music information retrieval and computational musicology have explored structured and multimodal understanding, few resources support factual and contextual music question answering (MQA) grounded in artist metadata or historical context. We introduce MusWikiDB, a vector database of 3.2M passages from 144K music-related Wikipedia pages, and ArtistMus, a benchmark of 1,000 questions on 500 diverse artists with metadata such as genre, debut year, and topic. These resources enable systematic evaluation of retrieval augmented generation (RAG) for MQA. Experiments show that RAG markedly improves factual accuracy—open-source models gain up to +56.8 percentage points (pp; Qwen3 8B: 35.0→91.8), approaching proprietary performance. RAG-style fine-tuning further boosts both factual recall and contextual reasoning, yielding strong improvements on both in-domain and out-of-domain benchmarks. MusWikiDB also yields +6 pp higher accuracy and 40% faster retrieval than the general Wikipedia corpus. We release MusWikiDB and ArtistMus to advance research in music information retrieval and domain-specific QA, establishing a foundation for retrieval augmented reasoning in culturally rich domains such as music.

Keywords: Music Information Retrieval, Question Answering, Retrieval Augmented Generation, Benchmark Dataset, Knowledge Base

1. Introduction

Large Language Models (LLMs) have demonstrated impressive capabilities across diverse domains, yet their effectiveness in music-related tasks remains limited due to sparse music-specific coverage in pretraining data. This limitation is particularly evident in Music Question Answering (MQA), where users seek detailed information about artists, genres, instruments, and music history—knowledge that general-purpose LLMs often lack or hallucinate (Ramoneda et al., 2024).

Traditional approaches to domain adaptation rely on fine-tuning LLMs with domain-specific data (Jeong, 2024; Sahoo et al., 2024). However, this faces significant challenges: securing high-quality training data, escalating computational costs, and the persistent difficulty of continuously updating model knowledge as new music emerges. Retrieval Augmented Generation (RAG) (Lewis et al., 2021) offers a compelling alternative by grounding model outputs in external knowledge bases, enabling dynamic access to specialized information without retraining.

Consider a user querying: *"Which jazz artists influenced Thad Jones's move to Copenhagen?"* General-purpose LLMs often fail at such queries, either hallucinating biographical details or providing generic responses lacking specific career contexts. While fine-tuning could address this, it requires ex-

pensive retraining for each knowledge update (e.g., new album releases, career milestones). Moreover, music knowledge is inherently dynamic and culturally situated, making static parametric storage inadequate. This motivates a retrieval-based approach where models can access up-to-date, verifiable information while maintaining the flexibility to reason over retrieved context.

Despite RAG's success in domains such as medicine (Miao et al., 2024) and law (Wiratunga et al., 2024), its application in music remains underexplored due to the absence of music-specific retrieval resources. Existing knowledge bases like Wikipedia lack the density and structure needed for effective MQA. Similarly, current music benchmarks (Weck et al., 2024; Yuan et al., 2024; Ramoneda et al., 2024) focus primarily on music theory or multimodal understanding, leaving a critical gap in evaluating artist-centric knowledge—information central to everyday music discovery and consumption (Lee, 2010; Doh et al., 2024).

To address these gaps, we introduce two foundational resources for music-domain question answering. First, MusWikiDB is a specialized knowledge base containing 3.2M passages from 144K music-related Wikipedia pages, covering seven categories: artists, genres, instruments, history, technology, theory, and forms. Unlike general Wikipedia corpora, MusWikiDB is optimized for music retrieval with domain-specific indexing and passage segmentation. Second, ArtistMus is a benchmark

* Corresponding author.

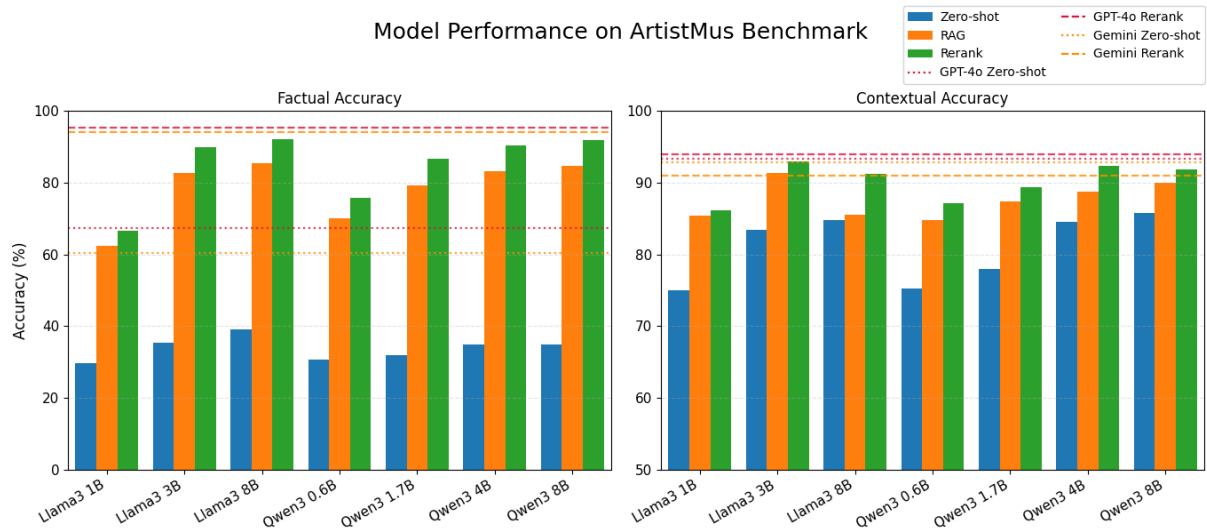


Figure 1: Performance comparison of open-source LLMs on the ArtistMus benchmark under three settings (Zero-shot, RAG, and Rerank). The dashed lines denote closed-source models (GPT-4o and Gemini 2.5 Flash) evaluated in the same conditions. The Rerank strategy consistently yields the highest factual and contextual accuracy across all open models, narrowing the gap with proprietary systems.

of 1,000 questions covering 500 globally diverse artists (spanning 163 countries), with rich metadata including genre, debut year, and topic category.

Using these resources, we conduct systematic experiments showing that retrieval-augmented approaches markedly improve music QA performance. Open-source models equipped with MusWikiDB-based retrieval achieve substantial gains in factual accuracy, performing on par with or exceeding proprietary systems. Moreover, RAG-style training on (*context, question, answer*) triples further enhances factual recall and contextual reasoning. The framework generalizes well to out-of-domain benchmarks and yields more accurate and efficient retrieval than the general Wikipedia corpus.¹

Our contributions are as follows:

- We release **MusWikiDB**, a comprehensive music-specific vector database containing 3.2M passages from 144K music-related Wikipedia pages, designed for retrieval-augmented music QA.
- We present **ArtistMus**, a benchmark of 1,000 validated questions on 500 globally diverse artists, with rich metadata enabling fine-grained analysis.
- Using MusWikiDB and ArtistMus, we apply RAG to the music domain, achieving substantial improvements in factual accuracy and en-

abling open-source models to approach proprietary performance.

2. Related Works

2.1. Music Question Answering Benchmarks

Question Answering (QA) is a core task in Information Retrieval, involving the provision of accurate answers to given questions by retrieving relevant information from document collections (Allam and Haggag, 2012). While open-domain QA systems operate over broad knowledge bases (Rajpurkar et al., 2016; Karpukhin et al., 2020), domain-specific QA targets specialized fields such as medicine (Pal et al., 2022), law (Chalkidis et al., 2021), or music, where both documents and queries are confined to domain-specific knowledge.

Several recent benchmarks have emerged to evaluate LLM performance on music-related tasks. MuChoMusic (Weck et al., 2024) presents 1,187 audio-based multiple-choice questions assessing musical knowledge and reasoning. MusicTheoryBench (Yuan et al., 2024) contains 372 expert-validated questions evaluating advanced music theory knowledge. TrustMus (Ramoneda et al., 2024) comprises 400 questions across four domains—People, Instruments & Technology, Genres/Forms/Theory, and Culture & History—derived from The Grove Dictionary Online (Sadie and Tyrrell, 2001). ZIQI-Eval (Li et al., 2024) offers 14,000 comprehension tasks spanning 10 major topics and 56 subtopics.

¹Code and data at <https://github.com/DaeyongKwon98/ArtistMus>

However, a critical gap exists in existing benchmarks: the inadequate representation of artist-centric metadata—information crucial for everyday music listening contexts (Lee, 2010; Doh et al., 2024). Current text-only QA benchmarks lack comprehensive coverage of details that listeners frequently seek, such as complete discographies, collaboration networks, creative evolution across albums, and career achievements. This disparity between current MQA capabilities and practical information demands motivates the development of artist-focused evaluation resources.

2.2. Retrieval Augmented Generation for Domain Adaptation

Domain adaptation of LLMs traditionally relies on fine-tuning with domain-specific data (Jeong, 2024; Sahoo et al., 2024; Satterfield et al., 2024). However, this approach faces challenges: securing high-quality training data, escalating computational costs with model scale, and difficulties in continuously updating model knowledge.

Retrieval Augmented Generation (RAG) (Lewis et al., 2021) offers an alternative by augmenting LLMs with external knowledge retrieval. Rather than relying solely on parametric memory, RAG retrieves relevant passages from external databases during inference to ground responses in factual context. Prior work demonstrates that RAG improves factual accuracy, reduces hallucination, enables transparent source verification, and provides greater scalability and cost-efficiency compared to fine-tuning alone (Borgeaud et al., 2022; Yasunaga et al., 2022; Wang et al., 2023).

While RAG has been successfully applied in general-purpose (Ram et al., 2023) and domain-specific settings such as law (Wiratunga et al., 2024) and medicine (Miao et al., 2024), its application in the music domain remains underexplored. Music QA requires both factual recall (e.g., artist names, compositions, release dates) and contextual reasoning (e.g., cultural influences, historical relations), making it an ideal testbed for evaluating RAG’s effectiveness in culturally rich, knowledge-intensive domains.

Recent work has also explored RAG-style fine-tuning, where models are trained on (context, question, answer) triples rather than traditional (question, answer) pairs (Shuster et al., 2021). This approach explicitly teaches models to integrate supporting passages, potentially improving their ability to leverage external evidence. However, systematic evaluation of this strategy within specialized cultural domains like music has been limited.

Our work addresses these gaps by introducing music-specific resources for RAG-based MQA and providing empirical evaluation of both retrieval-

based inference and RAG-style fine-tuning in this domain.

3. Dataset

Our resource design follows three key principles to ensure long-term utility for the research community. First, reproducibility: we document all collection and processing steps to enable replication and extension to other music knowledge sources. Second, diversity: both MusWikiDB and ArtistMus prioritize global representation, moving beyond Western-centric coverage typical of existing resources. Third, rich annotation: metadata such as genre, debut year, and topic enable fine-grained analysis.

3.1. MusWikiDB

To address the lack of a music-specific vector database for RAG in MQA, we developed MusWikiDB. We began by collecting music-related content from Wikipedia across seven categories: *artists*, *genres*, *instruments*, *history*, *technology*, *theory*, and *forms*. These categories were selected to cover a broad spectrum of music knowledge, providing a well-rounded foundation for answering music-related questions. The data collection process began with seven representative Wikipedia pages as entry points. From each entry page, we followed the hyperlinks embedded within the page to gather the first layer of linked pages, denoted as depth 1. Next, for every page at depth 1, we further collected the pages linked within them, which we refer to as depth 2. Finally, the same procedure was applied once more to the depth 2 pages, retrieving their linked pages as depth 3. In this way, the crawling process expanded from the initial seven root pages outward, capturing increasingly detailed subtopics and related information up to a depth of 3. We split the content into wikipedia sections such as *history*, *background*, and *early life*. We then removed sections shorter than 60 tokens to ensure the remaining text had enough context for meaningful retrieval. Table 1 shows an example of the source data of MusWikiDB (MusWikiDB_raw). Table 2 summarizes the number of Wikipedia pages and the average token count across the seven categories.

To construct the RAG database from MusWikiDB_raw, the text was segmented into passages of up to 256 tokens with a 10% overlap between adjacent segments, based on the ablation study in Section 4.3, in order to preserve contextual continuity. For first-stage retrieval, we employed BM25 (Robertson and Walker, 1994), a classical yet highly effective algorithm for text relevance ranking, to build an efficient index for

Category	Page	Section	Text
Instruments	Wind instrument	Families	Modern brass instruments generally come in ...
Genres	Music of Tanzania	Singeli	Singeli is a music genre that originated in ...

Table 1: Examples from MusWikiDB_raw. It includes the category, page title, section name, and corresponding text.

Category	Page Count	Avg. Tokens
Music History	40,692	312
Artist	31,847	528
Genres	28,553	324
Instruments	27,490	316
Musical Forms	7,681	331
Music Theory	6,055	356
Music Technology	2,071	301
Total	144,389	331

Table 2: Number of unique Wikipedia pages and average token count per category in MusWikiDB_raw.

	MusWikiDB	Wikipedia Corpus
# Pages	0.14M	3.2M
# Passages	3.2M	21M
Total tokens	0.36B	2.1B
Vocab Size	7.5M	21.5M

Table 3: MusWikiDB and Wikipedia Corpus (Karpukhin et al., 2020) statistics.

MusWikiDB. This enabled fast retrieval of relevant information during RAG-based inference, thereby improving both the accuracy and contextual relevance of the generated answers. The resulting MusWikiDB serves as a scalable and up-to-date knowledge base, enhancing RAG performance in MQA tasks and enabling the system to address complex, domain-specific music-related questions with greater precision and context.

Table 3 compares our proposed MusWikiDB with the Wikipedia corpus (Karpukhin et al., 2020). While MusWikiDB contains fewer pages (0.14M vs 3.2M) and has a smaller vocabulary size (7.5M vs 21.5M), it consists exclusively of music-specialized text information.

3.2. ArtistMus

While previous benchmarks have mainly addressed general music understanding, none has focused on music metadata, particularly the artist, which is crucial in music listening contexts (Lee, 2010; Doh et al., 2024). Therefore, we introduce ArtistMus, a multiple-choice benchmark dataset on music artists, to evaluate the performance of LLMs

Example Questions of ArtistMus

Metadata:

Artist: Robyn
Country: Sweden
Debut Year: 1989
Genre: Pop
Topic: Artistry

Factual Question:

"What vocal range does Robyn have?"

- A. Alto
- B. Soprano
- C. Tenor
- D. Baritone

Answer: B

Contextual Question:

"How did Robyn's approach to music change while writing her eighth studio album, Honey, according to the text?"

- A. Robyn explored music that was hypnotic and lacked a traditional structure.
- B. Robyn aimed to replicate the success of her previous hits.
- C. Robyn focused on creating more mainstream pop hits.
- D. Robyn returned to her R&B roots with a focus on tidy pop songs.

Answer: A

in artist-related QA using artist-specific data from MusWikiDB.

3.2.1. Data Collection

We selected music artists from MusWikiDB_raw whose infoboxes contained both *genre* and *year* information. The sections from their Wikipedia pages were then ranked by frequency, and the top five topics—*biography*, *career*, *discography*, *artistry*, and *collaborations*—were chosen to capture the broadest range of artist-related content. To balance efficiency and contextual richness, we retained text passages with token lengths from 500 to 2000.

For genre normalization (Schreiber, 2015), all genre labels were converted to lowercase, and spaces, hyphens (-), and slashes (/) were removed.

We used the 48 root genres defined in (Schreiber, 2016) as our initial genre taxonomy. After filtering to retain data corresponding to the 300 most frequent genres, each was mapped to one of 20 final genre labels.

To extract artists’ regional information, we provided the first paragraph (abstract) of each Wikipedia page to Llama 3.1 8B Instruct (Grattafiori et al., 2024), prompting it to identify the country associated with each artist. The list of valid countries was obtained from the *pycountry*² library.

3.2.2. Data Selection

We selected a diverse set of 500 artists based on their *topic*, *genre*, *debut year*, and *country*. Among these factors, *country* was given the highest priority, with a focus on including artists from underrepresented regions. Specifically, we applied an inverse frequency strategy on MusWikiDB, prioritizing countries with lower representation to ensure geographical balance. This approach enabled the inclusion of artists from minor regions such as Mauritius, Senegal, and Azerbaijan.

Figure 2 illustrates the regional distribution of the artists included in ArtistMus.³ Unlike many existing datasets that are heavily biased toward the U.S. and Europe, ArtistMus deliberately incorporates artists from 163 countries (or regions), ensuring broad global diversity and representation. Approximately 43% of the artists originate from outside the U.S. and Europe, indicating balanced representation beyond Western and European cultural centers.

To further enhance diversity and balance, we replaced dominant genres and topics with less common ones. Instead of focusing on widely represented genres such as *Pop* or *Rock*, we prioritized underrepresented genres such as *World Music* and *Latin*. Similarly, rather than overrepresented topics like *Biography*, we emphasized less frequent ones such as *Artistry*. In addition, debut years were selected to achieve a more uniform temporal distribution. The distribution of debut years for the artists in ArtistMus is shown in Figure 3.

For each artist, we generated one factual and one contextual question to evaluate the LLM’s factual accuracy and contextual reasoning ability. To construct these questions, we provided GPT-4o (Achiam et al., 2023) with the corresponding section text. Factual questions target verifiable details such as dates, names, and events, whereas contextual questions require reasoning or synthesis across multiple pieces of information within the passage.

²<https://pypi.org/project/pycountry/>

³You can see larger figure at Figure 5.




Figure 2: Regional distribution of the 500 music artists in ArtistMus, spanning 163 countries (or regions) to ensure global diversity beyond the traditional U.S.- and Europe-centric focus.




Figure 3: Debut years of the 500 music artists in ArtistMus.

3.2.3. Data Validation

To ensure benchmark quality, we applied a two-stage validation process. We validate the generated questions based on two criteria: *Music Relevance* and *Faithfulness*. For Music Relevance, questions that did not pertain to musical aspects were excluded except important details such as the artist’s birthplace. For Faithfulness, we prompted the LLM to verify whether the question and answer could be derived from the provided text. Finally, 1,000 multiple-choice questions passing final human validation were generated. We adjusted the placement of correct answers to ensure an even distribution, assigning 250 instances to each option.

In addition to the question–answer pairs, ArtistMus provides richer annotations, including the artist identity, topic category, the supporting text passage, as well as metadata such as genre and debut year. This richer structure enables fine-grained evaluation, for example analyzing performance across genres, eras, or topical categories.

4. Experiments

4.1. Benchmarks

For evaluation, we used two datasets: ArtistMus (in-domain) and TrustMus (out-of-domain) (Ramoneda

et al., 2024).⁴ Performance on factual and contextual questions was separately measured on the ArtistMus. For TrustMus, evaluation was conducted across four categories: People, Instrument & Technology, Genre, Forms, and Theory, and Culture & History, each comprising 100 questions. All evaluations use a multiple-choice QA format, and accuracy (exact match of the predicted option A, B, C, or D) was adopted as the evaluation metric.

4.2. Models and Experimental Setup

We evaluate MQA performance across two complementary experimental settings. All experiments were conducted in a deterministic setting with temperature set to 0.

Main Experiments. Main experiments measure zero-shot, RAG, and Rerank performance across a diverse set of LLMs. Specifically, we evaluate:

- **Zero-shot:** GPT-4o (Achiam et al., 2023) (API-based), Gemini 2.5 Flash (API-based), Chat-Musician (Yuan et al., 2024) (music-specific), Llama 3 (1B/3B/8B) (Grattafiori et al., 2024), and Qwen3 (0.6B/1.7B/4B/8B) (Yang et al., 2025). For Qwen3 models, we disabled the thinking mode to ensure consistent behavior.
- **RAG:** Each model is equipped with first-stage retrieval using BM25 over MusWikiDB, using 256-token passages as identified as optimal in Section 4.3.
- **Rerank:** We add a second-stage reranking step with the BGE reranker (Xiao et al., 2023), selecting the most relevant passages for generation.

This setup enables a systematic comparison across both closed-source and open-source models of varying scales, as well as a music-specific model, under different retrieval strategies.

Ablation Study. To disentangle the impact of adaptation strategies, we further compare three training paradigms on a fixed base model (Llama 3.1 8B Instruct):

- **Zero-shot baseline:** No adaptation.
- **QA fine-tuning:** Fine-tuned on 8K multiple-choice QA pairs from MusWikiDB.
- **RAG-style fine-tuning:** Fine-tuned on 8K augmented (*context, question, answer*) triples, enabling the model to leverage external passages during training.

⁴Performance on the MusicTheoryBench (Yuan et al., 2024) and ZIQI-Eval (Li et al., 2024) benchmarks will be included in the camera-ready version.

This ablation isolates the effects of standard QA fine-tuning versus RAG-style fine-tuning, complementing the experiment by highlighting the benefits of incorporating contextual grounding into training.

We adopt the following training configuration for both QA fine-tuning and RAG-style fine-tuning: All fine-tuned models were trained for one epoch using LoRA (Hu et al., 2022) with 8-bit quantization. We set the following hyperparameters: batch size = 2, gradient accumulation steps = 4, learning rate = $3e-5$, weight decay = 0.005, warmup ratio = 0.1, cosine learning rate scheduler (Wolf et al., 2020), AdamW optimizer (Loshchilov and Hutter, 2017), LoRA rank $r = 16$, scaling factor $\alpha = 16$, and dropout = 0.1.

4.3. Retriever Configurations

We conducted an additional study on the ArtistMus benchmark to determine the optimal retriever configuration for MusWikiDB as followings:

- **Passage size:** 128, 256, or 512 tokens.
- **Embedding models:** BM25 (Robertson and Walker, 1994) (sparse), Contriever (Izacard et al., 2021) (dense).
- **Reranker:** bge-reranker-large (Xiao et al., 2023).

For a fair comparison, we fixed the total token budget to 1024 by adjusting the number of retrieved passages: top-8 for 128-token passages, top-4 for 256, and top-2 for 512.

5. Results

5.1. Main Experiments: Zero-shot vs. RAG vs. Reranker

In-domain (ArtistMus). Table 4 summarizes the results of a wide range of LLMs under three inference settings: zero-shot, RAG, and RAG+Reranker. Across all models, factual questions were notably more challenging than contextual ones. For example, GPT-4o outperformed Llama 3.1 8B by 28.4 pp on factual questions in zero-shot setting, while the gap narrowed to 8.2 pp on contextual questions. RAG substantially improved factual accuracy across open-source models: Llama 3.1 3B improved from 39.0% to 85.4% (+46.4 pp), and Qwen3 8B from 35.0% to 84.8% (+49.8 pp). Adding a reranker further boosted accuracy, with gains of up to +7.6 pp compared to RAG. These results confirm that retrieval is more beneficial for entity-centric questions such as artist names and works.

GPT-4o and Gemini 2.5 Flash achieved the strongest overall performance, but even smaller open-source models (e.g., Llama 3.1 3B, Qwen3

			Zero-shot		RAG		Rerank	
	Model	Params	Fact.	Context.	Fact.	Context.	Fact.	Context.
Closed-source	GPT-4o	N/A	67.4	93.0	93.0	93.0	95.4	94.0
	Gemini 2.5 Flash	N/A	60.4	93.2	86.6	90.0	94.2	91.0
Music-specific	ChatMusician	7B	26.8	72.2	32.4	39.4	36.4	47.4
Open-source	Llama 3	1B	29.6	75.0	62.4	85.4	66.6	86.2
		3B	35.4	83.4	82.8	91.4	90.0	93.0
		8B	39.0	84.8	85.4	85.6	92.2	91.2
	Qwen 3	0.6B	30.6	75.2	70.0	84.8	75.8	87.2
		1.7B	32.0	78.0	79.2	87.4	86.8	89.4
		4B	34.8	84.6	83.2	88.8	90.4	92.4
		8B	35.0	85.8	84.8	90.0	91.8	91.8

Table 4: Performance on the ArtistMus benchmark across models and retrieval strategies. RAG uses BM25-based retrieval (top-4 passages, 256 tokens). Rerank applies BGE reranker to refine passage selection. Factual questions (Fact.) test entity recall (e.g., names, dates); contextual questions (Context.) require reasoning over narrative information. All results are accuracy (%) on 500 questions per type.

Models	# Params	Zero-shot / RAG / Rerank
Llama 3	1B	0.25 / 0.36 / 0.38 (+13%)
	3B	0.29 / 0.38 / 0.39 (+10%)
	8B	0.36 / 0.41 / 0.43 (+7%)
Qwen 3	0.6B	0.28 / 0.36 / 0.38 (+10%)
	1.7B	0.33 / 0.39 / 0.39 (+6%)
	4B	0.32 / 0.38 / 0.41 (+9%)
	8B	0.34 / 0.40 / 0.43 (+9%)

Table 5: Performance on the TrustMus benchmark. In this setting, RAG corresponds to first-stage retrieval based on BM25, whereas Rerank represents two-stage retrieval employing the BGE reranker.

1.7B) benefited substantially from RAG and reranking. With retrieval support, these smaller models approached the performance of much larger API-based models, demonstrating that RAG can effectively bridge the gap between open-source and proprietary systems. ChatMusician, despite being designed for music-specific tasks, achieved only around 26% accuracy—comparable to random guessing—largely due to frequent instruction-following failures. Similarly, the RAG setting exhibited degraded contextual accuracy, primarily because the models often failed to properly follow task-specific instructions when integrating retrieved information.

Out-of-domain (TrustMus). Table 5 presents results on the TrustMus benchmark, which represents a different knowledge source than Wikipedia. Despite this shift, our methodology continued to provide consistent gains: for instance, Llama 3.1 1B improved from 25.0% to 36.0% (+11.0 pp) with RAG, and further to 38.0% (+2.0 pp) with reranking. This indicates that retrieval with MusWikiDB not only strengthens in-domain factual recall but also

transfers effectively to out-of-domain knowledge, even when the benchmark corpus differs from the retrieval database.

Analysis. The substantial factual accuracy gains from RAG (+40.4 pp on average) contrast with modest contextual improvements (+4.8 pp on average), revealing an important distinction⁵. Factual questions often require precise entity recall (e.g., album names) that are absent from model parameters but directly retrievable from passages. Contextual questions, which demand reasoning over information, benefit less from retrieval alone—models already possess adequate reasoning capabilities once factual grounding is provided. This suggests a division of labor: retrieval provides facts, while model parameters handle reasoning. The smaller gap between model sizes with RAG (e.g., Qwen 0.6B vs 8B: 77.4% vs 87.4% on average) further indicates that retrieval mitigates the need for large parametric memory, democratizing access to factual knowledge.

5.2. Ablation Study: Fine-tuning Strategies

Table 6 reports the results of different fine-tuning strategies on ArtistMus using the Llama 3.1 8B Instruct model. Several observations emerge:

- **QA Fine-tuning:** Compared to the zero-shot baseline (39.0% / 84.8%), QA fine-tuning slightly improved factual accuracy to 40.8% (+1.8 pp) but led to a large drop in contextual performance to 75.6% (−9.2 pp). This suggests that conventional QA fine-tuning may

⁵ChatMusician is excluded because of low instruction following rate.

	Factual	Contextual	Average
Zero-shot	39.0	84.8	61.9
QA FT	40.8	75.6	58.2
RAG	85.4	85.6	85.5
Rerank	92.2	91.2	<u>91.7</u>
RAG FT	86.8	<u>93.0</u>	89.9
RAG FT Rerank	92.2	94.0	93.1

Table 6: Ablation study. QA FT refers to the zero-shot performance of the QA fine-tuned model, RAG and Rerank are taken from the main experiment, and RAG FT and RAG FT Rerank correspond to the performance of the RAG-style fine-tuned model with RAG and reranker applied. The best result is highlighted in **bold**, while the second best is underlined.

overfit to training data, improving memorization while harming general inference ability.

- **RAG Inference vs. Reranking:** RAG inference yielded a dramatic gain in factual accuracy (+46.4 pp), while maintaining contextual accuracy. Adding a reranker further boosted both factual (+6.8 pp) and contextual (+5.6 pp) performance, showing the importance of two-stage retrieval.
- **RAG-style Fine-tuning:** Incorporating retrieved passages during training improved contextual understanding, with accuracy rising to 93.0% (+7.4 pp), while factual accuracy remained strong (86.8%). With reranker, it achieved the best overall performance (93.1%), confirming that training the model to leverage retrieved evidence enhances both factual recall and contextual reasoning.

5.3. Retriever Configuration Study

Table 7 shows the effect of varying passage size, embedding models, and the usage of reranker.

- Contextual accuracy generally improved with shorter passages, while factual accuracy was more stable across passage sizes for BM25 but degraded significantly for Contriever.
- BM25 with a reranker consistently outperformed dense retrieval methods, reflecting the entity-centric nature of ArtistMus.
- The best performance was achieved with BM25 + reranker at 256 tokens, balancing accuracy and efficiency.

As shown in Figure 4, MusWikiDB achieves both higher retrieval accuracy (+6.0 pp absolute improvement) and significantly faster (40%) retrieval speed compared to the general Wikipedia corpus.

Embedding	Passage Size	Factual	Contextual
BM25	128	81.2	87.2
	256	84.4	85.6
	512	85.0	85.8
BM25^R	128	85.6	89.6
	256	92.2	<u>91.2</u>
	512	<u>91.0</u>	89.8
Contriever	128	80.8	80.0
	256	69.2	84.6
	512	54.6	75.4
Contriever^R	128	86.8	82.0
	256	84.2	91.6
	512	70.6	87.6

Table 7: Llama 3.1 8B Instruct (Grattafiori et al., 2024) RAG performance on ArtistMus, by different passage size and embeddings. ^R indicates the use of a reranker. The best result is highlighted in **bold**, while the second best is underlined.

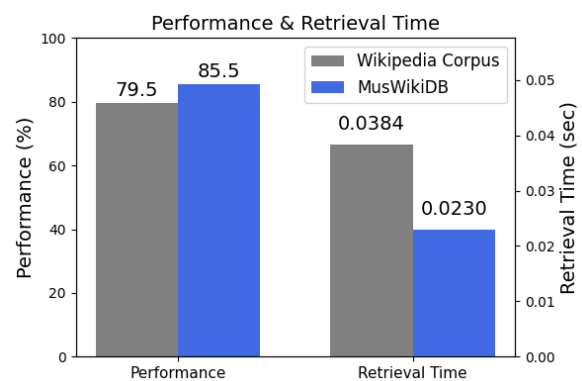


Figure 4: RAG performance and retrieval time for Wikipedia Corpus (Karpukhin et al., 2020) and MusWikiDB.

These results highlight the effectiveness of a music-specialized knowledge base: not only does it provide more relevant evidence for music question answering, but it also enables efficient large-scale inference, which is crucial for real-time applications.

6. Conclusion

This paper introduces two key contributions for music information retrieval: MusWikiDB, a domain-specific knowledge base with over 3M passages from 144K music-related Wikipedia pages, and ArtistMus, a benchmark of 1000 questions covering 500 globally diverse artists. Together, they fill critical gaps in music question answering (MQA).

Our experiments show that retrieval augmented generation (RAG) substantially boosts factual accuracy and overall robustness in MQA, narrowing the gap between open-source and proprietary systems. Notably, smaller models equipped with retrieval can match or even surpass large closed models,

demonstrating that domain-specific retrieval democratizes access to high-quality, knowledge-grounded MQA. MusWikiDB also enables more accurate and efficient retrieval than the general Wikipedia corpus, validating the value of domain-specialized corpora for real-time applications.

Furthermore, integrating retrieval into model training enhances both factual and contextual reasoning, establishing a new paradigm for unifying retrieval and generation in music-domain LLMs. The consistent gains observed across in-domain and out-of-domain benchmarks highlight the transferability of this approach.

7. Ethics Statement

MusWikiDB is constructed from Wikipedia, which, while comprehensive, reflects the biases inherent in crowdsourced knowledge bases, including underrepresentation of non-Western music traditions and artists from marginalized communities. While ArtistMus includes artists from 163 countries, the quality and depth of available information varies significantly by region and genre.

Although RAG substantially improves factual grounding, our system may still generate incorrect information, particularly for less-documented artists or culturally contested interpretations. We emphasize that generated outputs should be treated as starting points for exploration rather than authoritative sources, especially in applications where misinformation could impact artist reputations or cultural understanding.

This work is designed to advance research in music information retrieval and domain-specific QA systems. We caution against deploying such systems in high-stakes applications without human oversight and appropriate disclaimers. The benchmark and resources are released to enable further research into addressing these challenges.

8. Limitations

ArtistMus focuses on artist-centric questions with a multiple-choice format. While this enables consistent automatic evaluation, it does not capture all aspects of music question answering, such as open-ended explanations, comparative analysis across artists, or questions requiring multi-hop reasoning. The 1,000-question scale, though sufficient for establishing trends, could be expanded to enable more fine-grained analysis across subgenres and eras.

MusWikiDB relies exclusively on Wikipedia content, which provides broad coverage but may lack depth for niche genres, regional music traditions, or recently emerging artists. The crawling strategy based on hyperlink depth may also introduce

coverage gaps. Future work could integrate additional structured music databases to provide more comprehensive and up-to-date information.

9. Bibliographical References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ali Mohamed Nabil Allam and Mohamed Hassan Haggag. 2012. The question answering systems: A survey. *International Journal of Research and Reviews in Information Sciences (IJRRIS)*, 2(3).
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. 2022. Improving language models by retrieving from trillions of tokens. In *International conference on machine learning*, pages 2206–2240. PMLR.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Martin Katz, and Nikolaos Aletras. 2021. Lexglue: A benchmark dataset for legal language understanding in english. *arXiv preprint arXiv:2110.00976*.
- SeungHeon Doh, Keunwoo Choi, Daeyong Kwon, Taesu Kim, and Juhan Nam. 2024. Music discovery dialogue generation using human intent analysis and large language models. *arXiv preprint arXiv:2411.07439*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Unsupervised dense information retrieval with contrastive learning. *arXiv preprint arXiv:2112.09118*.
- Cheonsu Jeong. 2024. Fine-tuning and utilization methods of domain-specific llms. *arXiv preprint arXiv:2401.02981*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick SH Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *EMNLP (1)*, pages 6769–6781.
- Jin Ha Lee. 2010. Analysis of user needs and information features in natural language queries seeking music information. *Journal of the American Society for Information Science and Technology*, 61(5):1025–1045.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#).
- Jiajia Li, Lu Yang, Mingni Tang, Cong Chen, Zuchao Li, Ping Wang, and Hai Zhao. 2024. The music maestro or the musically challenged, a massive music evaluation benchmark for large language models. *arXiv preprint arXiv:2406.15885*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Jing Miao, Charat Thongprayoon, Supawadee Supadungsuk, Oscar A Garcia Valencia, and Wisit Cheungpasitporn. 2024. Integrating retrieval-augmented generation with large language models in nephrology: advancing practical applications. *Medicina*, 60(3):445.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikanan Sankarasubbu. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, pages 248–260. PMLR.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-context retrieval-augmented language models. *Transactions of the Association for Computational Linguistics*, 11:1316–1331.
- Pedro Ramoneda, Emilia Parada-Cabaleiro, Benno Weck, and Xavier Serra. 2024. [The role of large language models in musicology: Are we ready to trust the machines?](#)
- Stephen E Robertson and Steve Walker. 1994. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval.

- In *SIGIR'94: Proceedings of the Seventeenth Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval, organised by Dublin City University*, pages 232–241. Springer.
- Stanley Sadie and John Tyrrell. 2001. *The New Grove Dictionary of Music and Musicians, 2nd edition*. Macmillan Publishers, London. Accessed 05-05-2024.
- Satya S Sahoo, Joseph M Plasek, Hua Xu, Özlem Uzuner, Trevor Cohen, Meliha Yetisgen, Hongfang Liu, Stéphane Meystre, and Yanshan Wang. 2024. Large language models for biomedicine: foundations, opportunities, challenges, and best practices. *Journal of the American Medical Informatics Association*, 31(9):2114–2124.
- Nolan Satterfield, Parker Holbrooka, and Thomas Wilcoxa. 2024. Fine-tuning llama with case law data to improve legal domain performance. *OSF Preprints*.
- Hendrik Schreiber. 2015. Improving genre annotations for the million song dataset. In *ISMIR*, pages 241–247.
- Hendrik Schreiber. 2016. Genre ontology learning: Comparing curated with crowd-sourced ontologies. In *ISMIR*, pages 400–406.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*.
- Yile Wang, Peng Li, Maosong Sun, and Yang Liu. 2023. Self-knowledge guided retrieval augmentation for large language models. *arXiv preprint arXiv:2310.05002*.
- Benno Weck, Ilaria Manco, Emmanouil Benetos, Elio Quinton, George Fazekas, and Dmitry Bogdanov. 2024. Muchomusic: Evaluating music understanding in multimodal audio-language models. *arXiv preprint arXiv:2408.01337*.
- Nirmalie Wiratunga, Ramitha Abeyratne, Lasal Jayawardena, Kyle Martin, Stewart Massie, Ikechukwu Nkisi-Orji, Ruvan Weerasinghe, Anne Liret, and Bruno Fleisch. 2024. Cbr-rag: case-based reasoning for retrieval augmented generation in llms for legal question answering. In *International Conference on Case-Based Reasoning*, pages 445–460. Springer.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. *Transformers: State-of-the-art natural language processing*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. *C-pack: Packaged resources to advance general chinese embedding*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Michihiro Yasunaga, Antoine Bosselut, Hongyu Ren, Xikun Zhang, Christopher D Manning, Percy S Liang, and Jure Leskovec. 2022. Deep bidirectional language-knowledge graph pretraining. *Advances in Neural Information Processing Systems*, 35:37309–37323.
- Ruibin Yuan, Hanfeng Lin, Yi Wang, Zeyue Tian, Shangda Wu, Tianhao Shen, Ge Zhang, Yuhang Wu, Cong Liu, Ziya Zhou, et al. 2024. Chatmusician: Understanding and generating music intrinsically with llm. *arXiv preprint arXiv:2402.16153*.

A. Appendix

A.1. Examples of TrustMus

Category: People

Question:

"When did Adolph Martin Schlesinger found the music-publishing house?"

- A. In 1819
- B. Before 1795
- C. In August 1814
- D. In April 1810

Answer: D

Category: Instrument & Technology

Question:

"What is the name of the highland bagpipe that was brought by Scottish regiments during the British Raj?"

- A. Tablkhāna
- B. Bagpipe
- C. Naqqārakhāna
- D. Nobat

Answer: B

Category: Culture & History

Question:

"When did Amalarius of Metz die?"

- A. c830
- B. c825
- C. c850
- D. c775

Answer: C

Category: Genre & Forms & Theory

Question:

"What did Christoph Bernhard describe as 'harmonic counterpoint' in the 17th century?"

- A. The regulated alternation of perfect and imperfect consonances
- B. The relationships between such sounds
- C. The simultaneous sounding of two or three notes in isolation
- D. An articulated sequence of 'well-juxtaposed consonances and dissonances'

Answer: D

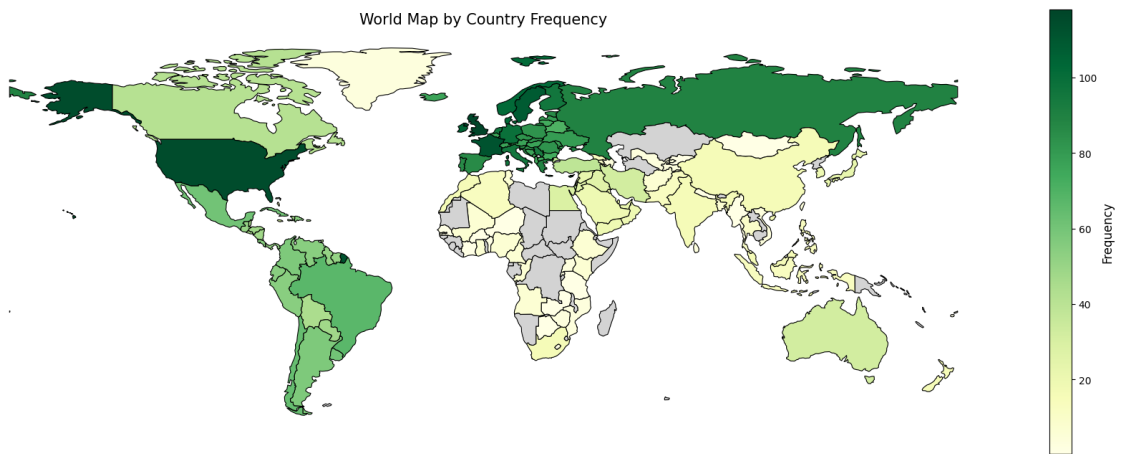


Figure 5: Regional distribution of the 500 music artists in ArtistMus, spanning 163 countries (or regions) to ensure global diversity beyond the traditional U.S.- and Europe-centric focus.