

Transformer Tafsir at QIAS 2025 Shared Task: Hybrid Retrieval-Augmented Generation for Islamic Knowledge Question Answering

Muhammad Abu Ahmad¹, Mohamad Ballout¹, Raia Abu Ahmad², Elia Bruni¹

¹Institute of Cognitive Science, University of Osnabrück, Osnabrück, Germany

²Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI), Berlin, Germany

Corresponding author: mabuahmad@uni-osnabrueck.de

Abstract

This paper presents our submission to the QIAS 2025 shared task on Islamic knowledge understanding and reasoning. We developed a hybrid retrieval-augmented generation (RAG) system that combines sparse and dense retrieval methods with cross-encoder reranking to improve large language model (LLM) performance. Our three-stage pipeline incorporates BM25 for initial retrieval, a dense embedding retrieval model for semantic matching, and cross-encoder reranking for precise content retrieval. We evaluate our approach on both subtasks using two LLMs, Fanar and Mistral, demonstrating that the proposed RAG pipeline enhances performance across both, with accuracy improvements up to 25%, depending on the task and model configuration. Our best configuration is achieved with Fanar, yielding accuracy scores of 45% in Subtask 1 and 80% in Subtask 2.

1 Introduction

QIAS 2025 is a question answering (QA) shared task that aims to evaluate large language models' (LLMs) ability to understand and reason within Islamic knowledge (Boucekif et al., 2025a,b). The task is divided into two subtasks: (1) Islamic Inheritance Reasoning, requiring precise application of inheritance law principles, and (2) Islamic Assessment, covering general Islamic knowledge across different topics such as theology, jurisprudence, biography, and ethics. Islamic jurisprudence (Fiqh) and inheritance law ('Ilm al-Mawārīth) present unique challenges within natural language processing (NLP) in the Arabic language, as it highlights the differences between Modern Standard Arabic (MSA) and Classical Arabic used in religion-related texts.

We tackle the QIAS shared task using a hybrid, naive retrieval-augmented generation (RAG) pipeline specifically designed for Arabic Islamic

knowledge. As shown in Figure 1, our approach consists of four components: 1. Preprocessing; 2. Three-stage hybrid retrieval pipeline; 3. Context integration; and 4. LLM inference. Our suggested retrieval system combines sparse retrieval via BM25 (Robertson et al., 2009), dense retrieval using Arabic-optimized embeddings (Nacar et al., 2025), and a miniLM-based cross-encoder reranking model for final passage selection.

In this paper, we present our system in detail and discuss our main contributions, including a specialized Arabic preprocessing pipeline, a hybrid three-stage retrieval architecture that combines complementary retrieval methods, a comprehensive evaluation across multiple LLMs demonstrating consistent improvements RAG context integration, and an analysis of the pipeline's performance. To facilitate reproducibility, we make our implementation publicly available.¹

2 Background

Task Setup. The QIAS 2025 shared task features two scholar-verified multiple-choice question (MCQ) subtasks with three difficulty levels. Subtask 1 (Islamic Inheritance Reasoning) covers 'Ilm al-Mawārīth with 9,450 training, 1,500 validation, and 1,000 test questions, plus 32,000 fatwas. Subtask 2 (Islamic General Knowledge) tests general Islamic knowledge from a selection of 25 classical Islamic knowledge books with 800 validation and 1,000 test questions. The complete corpus of books was also provided by the organizers.

Related Work. QA tasks have demonstrated significant benefits from retrieval-augmented pipelines, particularly when domain-specific knowledge bases exist (Arslan et al., 2024). RAG combines retrieval with generative models by first retrieving relevant passages from a knowledge base,

¹<https://gitlab.com/mhauesh/qias-shared-task-2025-solution-implementation>

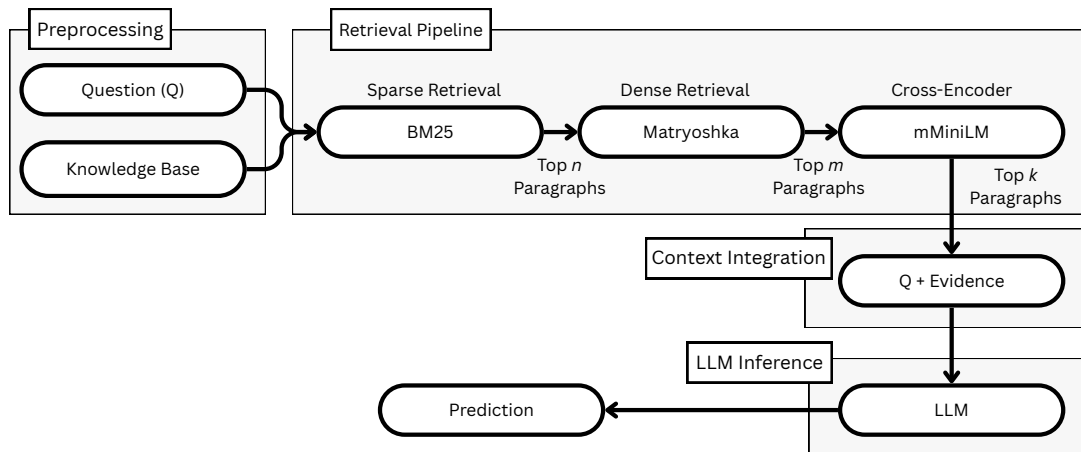


Figure 1: Proposed retrieval-augmented pipeline for the QIAS shared task.

then feeding them as context to language models for answer generation (Gao et al., 2023).

Modern hybrid retrieval pipelines typically consist of three stages: sparse retrieval, dense retrieval, and cross-encoder reranking (Huyen, 2024). BM25, a lexical retrieval method using TF-IDF and document-length normalization, serves as the most widely-used sparse retriever. Dense retrievers embed queries and passages into shared vector spaces using transformer models (Karpukhin et al., 2020), capturing semantic relationships that lexical methods might miss. Finally, cross-encoders provide higher precision by jointly scoring query-document pairs with full attention (Cheng et al., 2023).

Recent Arabic-focused RAG research demonstrates the value of this approach. For example, Arabica QA (Abdallah et al., 2024) presents QA pairs with a dense retrieval model pre-trained on Arabic for open-domain QA, proving the effectiveness of RAG pipelines on performances of various LLMs. Similarly, Al-Rasheed et al. (2025) show that integrating RAG pipelines improves LLM prediction results over retrieval-free setups.

Recent advancements in Arabic embedding models (Nacar et al., 2025), have significantly improved representation quality for retrieval tasks, with custom Matryoshka embeddings performing highly on the MTEB leaderboard.² These models offer compact yet powerful embeddings well-suited for scalable retrieval and reranking in Arabic-language RAG pipelines. In the context of Islamic and Quranic QA, however, additional challenges arise due to the linguistic divergence between MSA queries and Classical Arabic source

texts, semantic ambiguity in Quranic language, and the scarcity of high-quality, domain-specific datasets (Oshallah et al., 2025). Prior work has demonstrated the potential of retrieval-augmented QA systems (Khalila et al., 2025), showing that even small LLMs can generate relevant and faithful answers when grounded in appropriate retrieval context. Other studies have explored individual components, such as dense retrievers and rerankers for Arabic QA (Alsubhi et al., 2025; El-Beltagy and Abdallah, 2024).

Building on this foundation, our work contributes a task-specific, modular retrieval architecture that combines sparse, dense, and reranking components in a unified pipeline tailored for Classical Arabic QA in the domain of Islamic knowledge, providing insight into the utility of a hybrid retrieval system under realistic conditions.

3 System Overview

The architecture of the pipeline consists of four main components: Arabic text preprocessing and knowledge base construction, three-stage hybrid retrieval, context integration, and LLM inference.

Arabic Text Preprocessing. Arabic text processing poses unique challenges due to rich morphology, orthographic variations, and diacritical marks (Habash, 2010). Since each component in our retrieval pipeline requires specific preprocessing needs, we implement a two-tier preprocessing approach: *Full preprocessing* and *light preprocessing*. The former is used for BM25 indexing, and includes stopword removal using enhanced NLTK (Bird and Loper, 2004) lists, tokenization via CAMEL tools (Obeid et al., 2020), and token filtering by length and content criteria. On the

²<https://huggingface.co/spaces/mteb/leaderboard>

other hand, light preprocessing is used for dense retrieval and cross-encoder inputs and includes formatting normalization, punctuation removal, citation removal, character normalization, and dediacritization using CAMEL tools. The important distinguishing factor between the two preprocessing approaches is preserving semantic information when it comes to dense retrieval pipelines.

	Light	Full
Formatting	✓	✓
Punctuation removal	✓	✓
Dediacritization	✓	✓
Character normalization	✓	✓
Citations removal	✓	✓
Stopwords removal	✗	✓
Tokenization	✗	✓

Table 1: Preprocessing procedure for retrieval methods.

Knowledge Base Construction. We construct domain-specific knowledge bases from the provided training materials: For Subtask 1, we process 32,000 IslamWeb fatwas in JSON format. Each fatwa contains structured fields including category, question, answer, and metadata. We treat each complete fatwa as a single retrieval unit to maintain contextual coherence. For Subtask 2, we process classical Islamic books provided in HTML and DOCX formats, implementing paragraph-based chunking and applying overlap strategies to prevent information loss. We create tri-directional mappings between fully processed chunks, lightly processed chunks, and original text, enabling seamless integration across different retrieval stages.

Retrieval Pipeline. Our hybrid retrieval system combines complementary retrieval methods: 1. Sparse retrieval using BM25, which provides initial candidate selection using lexical matching using the fully preprocessed Arabic text, retrieving the top 1000 candidates. 2. Dense retrieval, applying semantic matching using Arabic-optimized embedding models. Based on existing benchmarks (Enevoldsen et al., 2025), we embed the lightly preprocessed text using Arabic-Triplet-Matryoshka-V2, selecting the top 200 passages closest to the question based on cosine similarity. 3. Cross-encoder reranking, which provides final precision enhancement using transformer models trained for relevance scoring. The cross-encoder jointly processes query-passage pairs, allowing full attention across inputs for more accurate relevance assessment. We use a miniLMv2 model fine-tuned

on the MMARCO dataset.³ Due to context window limitations in the used LLMs, we retrieve the top 5 passages with associated relevance scores for context integration in the LLM prompt.

Context Integration. We designed a prompt to support both tasks based on prior research (Schulhoff et al., 2024). First, we defined a domain-specific persona who is an expert Islamic scholar. Then, for each question, we retrieved relevant context passages and integrated them into the prompt using a format that prioritizes them as sources. We also added few-shot examples by selecting two random questions from the development set, demonstrating correct reasoning patterns, and included format constraints to enforce valid multiple-choice responses. We include the full prompt template in Appendix A.

4 Experimental Setup

We evaluate our proposed system across two distinct LLMs representing different model families and access patterns. First, we use Fanar (Team et al., 2025) via its API, which is a specialized Arabic LLM designed for Islamic content. Then, we experiment with Mistral (specifically, mistral-saba-24b),⁴ a state-of-the-art open-weight model, accessed through the Groq API.⁵ These LLMs were chosen since they were made available by the organizers of the QIAS shared task.

We use the same retrieval configuration on all tested LLMs: From BM25, we retrieve the top 1000 most relevant passages from the knowledge base, from dense retrieval, we filter those to the top 200 passages most similar to the given question, and finally, from the cross-encoder, we retrieve the top 5 passages to be integrated as context when prompting LLMs. We process the query and the knowledge base using CAMEL Tools (v1.2.0) for normalization, dediacritization, and tokenization. Additionally, we implement custom routines for citation removal, formatting cleanup, and chunking the knowledge base.

Performance of LLMs is measured using accuracy, which is the percentage of questions where the model’s prediction exactly matches the correct answer, evaluated during the testing phase of the shared task via the provided platform.

³<https://huggingface.co/cross-encoder/mmarco-miniLMv2-L12-H384-v1>

⁴<https://mistral.ai/news/mistral-saba>

⁵<https://console.groq.com/home>

5 Results & Discussion

We present our results in Table 2 for both subtasks, demonstrating varying levels of improvement when incorporating RAG across all tested configurations. The magnitude of improvement varies significantly between tasks and models. Subtask 1 (Islamic Inheritance) shows modest but consistent improvements of 1%-4% when implementing our proposed pipeline. We hypothesize that the specialized nature of inheritance law calculations may limit RAG effectiveness, as these problems often require precise mathematical reasoning rather than factual retrieval, as well as the high degree of semantic similarity between questions that require different methods to solve. However, we note that Subtask 2 (General Islamic Knowledge) exhibits substantial performance boosts with the RAG pipeline, with improvements ranging from 10%-25%. This dramatic enhancement suggests that general Islamic knowledge questions benefit significantly from access to authoritative source material.

Model Configuration	Accuracy Score
<i>Subtask 1: Islamic Inheritance</i>	
Fanar Baseline	44.0%
Fanar Transformer Tafsir	45.0%
Mistral Baseline	35.0%
Mistral Transformer Tafsir	39.0%
<i>Subtask 2: General Islamic Knowledge</i>	
Fanar Baseline	55.0%
Fanar Transformer Tafsir	80.0%
Mistral Baseline	69.0%
Mistral Transformer Tafsir	79.0%

Table 2: Results on the given test sets for both subtasks of the QIAS shared task using our proposed hybrid retrieval pipeline (+RAG) compared to baseline model performance (without RAG).

Comparing different LLMs, we note that Fanar shows smaller relative improvements (1%-25%) but reaches higher absolute performance on Subtask 2, likely due to its Islamic domain specialization and prior exposure to similar training data. On the other hand, Mistral demonstrates more consistent relative improvements (4%-10%) across both tasks, suggesting the overall benefit of RAG pipelines to improve performance of models on very specific domains, such as Islamic knowledge in this case, when similar in-domain data was most likely lacking in their training processes.

6 Error Analysis

A manual error analysis reveals three key patterns. First, models struggle with the complex, fractional reasoning in Subtask 1, indicating a need for symbolic reasoning beyond current RAG approaches. Second, errors arise when retrieved context is relevant but incomplete, highlighting the importance of comprehensive knowledge bases and accurate retrieval. Finally, some questions demand logical inference that simple retrieval cannot solve, suggesting a need for specialized training methodologies (Ke et al., 2025) or reasoning-based prompting (Qiao et al., 2023).

Analyzing by difficulty level for our best results (Fanar Transformer Tafsir for both tasks), we see varied performance by the LLMs on the two subtasks. Task 1 declined from 54.27% on beginner questions to 36.22% on advanced questions, while Subtask 2 showed a similar pattern (82.21% to 73.33%), as shown in Table 3. Subtask 2 errors were more evenly distributed, but showed difficulty in theological reasoning and jurisprudential methodology.

Task	Beginner	Intermediate	Advanced	Overall
Subtask 1	53.40%	–	36.00%	44.70%
Subtask 2	81.86%	78.67%	73.33%	80.10%

Table 3: System performance by task and difficulty level.

7 Conclusion

We presented a hybrid RAG pipeline for the QIAS 2025 shared task on Islamic knowledge QA. Our pipeline involved Arabic-specific preprocessing, a three-stage retrieval architecture (BM25, dense retrieval, cross-encoder reranking), context integration, and LLM inference. Based on evaluations of Fanar and Mistral, we showed that our method consistently outperformed baselines, demonstrating that RAG improves accuracy, especially for general knowledge over structured inheritance problems. To improve the presented system, future work can explore dynamic context selection, domain fine-tuning, and integrating structured reasoning modules for inheritance law.

Limitations

Our system struggles with complex inheritance problems requiring multi-step mathematical reasoning, as it lacks symbolic reasoning capabilities.

The presented system is not fine-tuned, hence the semantic similarity of recurring words and phrases in subtask 1 limits the system’s ability to retrieve precise relevant passages. The system is also dependent on the quality of the external knowledge base and does not explore knowledge curation. We did not use the provided training data, so domain fine-tuning remains unexplored. The number of retrieved passages in each step of the pipeline requires further investigation in order to fully maximize the system’s capabilities. Choosing said parameters (n,m and k) was based on trials conducted on the developments sets and not on the test sets. Finally, we did not explore performance on additional LLMs, which maybe have yielded better results, due to time and compute limitations.

Acknowledgments

We thank the QIAS 2025 organizers for providing comprehensive datasets and evaluation frameworks. We also acknowledge computational resources provided by the University of Osnabrück.

References

- Abdelrahman Abdallah, Mahmoud Kasem, Mahmoud Abdalla, Mohamed Mahmoud, Mohamed Elkasaby, Yasser Elbendary, and Adam Jatowt. 2024. Arabicaqa: A comprehensive dataset for arabic question answering. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2049–2059.
- Raghad Al-Rasheed, Abdullah Al Muaddi, Hawra Al-jasim, Rawan Al-Matham, Muneera Alhoshan, Asma Al Wazrah, and Abdulrahman AlOsaimy. 2025. *Evaluating RAG pipelines for Arabic lexical information retrieval: A comparative study of embedding and generation models*. In *Proceedings of the 1st Workshop on NLP for Languages Using Arabic Script*, pages 155–164, Abu Dhabi, UAE. Association for Computational Linguistics.
- Jumana Alsubhi, Mohammad D Alahmadi, Ahmed Alhusayni, Ibrahim Aldailami, Israa Hamdine, Ahmad Shabana, Yazeed Iskandar, and Suhayb Khayyat. 2025. *Optimizing rag pipelines for arabic: A systematic analysis of core components*. *arXiv preprint arXiv:2506.06339*.
- Muhammad Arslan, Hussam Ghanem, Saba Munawar, and Christophe Cruz. 2024. *A survey on rag with llms*. *Procedia computer science*, 246:3781–3790.
- Steven Bird and Edward Loper. 2004. *NLTK: The natural language toolkit*. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- Abdessalam Boucekif, Samer Rashwani, Mohammed Ghaly, Mutaz Al-Khatib, Emad Mohamed, Wajdi Zaghrouani, Heba Sbahi, Shahd Gaben, and Aiman Erbad. 2025a. *Qias 2025: Overview of the shared task on islamic inheritance reasoning and knowledge assessment*. In *Proceedings of The Third Arabic Natural Language Processing Conference, ArabicNLP 2025, Suzhou, China, November 5–9*. Association for Computational Linguistics.
- Abdessalam Boucekif, Samer Rashwani, Heba Sbahi, Shahd Gaben, Mutaz Al-Khatib, and Mohammed Ghaly. 2025b. *Assessing large language models on islamic legal reasoning: Evidence from inheritance law evaluation*. In *Proceedings of The Third Arabic Natural Language Processing Conference (ArabicNLP 2025)*, Suzhou, China. Association for Computational Linguistics.
- Hao Cheng, Hao Fang, Xiaodong Liu, and Jianfeng Gao. 2023. *Task-aware specialization for efficient and robust dense retrieval for open-domain question answering*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1864–1875, Toronto, Canada. Association for Computational Linguistics.
- Samhaa R El-Beltagy and Mohamed A Abdallah. 2024. *Exploring retrieval augmented generation in arabic*. *Procedia Computer Science*, 244:296–307.
- Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Siblini, Dominik Krzemiński, Genta Indra Winata, Saba Sturua, Saiteja Utpala, Mathieu Ciancone, Marion Schaeffer, Gabriel Sequeira, Diganta Misra, Shreeya Dhakal, Jonathan Ryrström, Roman Solomatin, and 67 others. 2025. *Mmteb: Massive multilingual text embedding benchmark*. *arXiv preprint arXiv:2502.13595*.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yixin Dai, Jiawei Sun, Haofen Wang, and Haofen Wang. 2023. *Retrieval-augmented generation for large language models: A survey*. *arXiv preprint arXiv:2312.10997*, 2(1).
- Nizar Y Habash. 2010. *Introduction to Arabic natural language processing*. Morgan & Claypool Publishers.
- Chip Huyen. 2024. *AI Engineering: Building Applications with Foundation Models*. O’Reilly Media, Incorporated.
- 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.
- Zixuan Ke, Fangkai Jiao, Yifei Ming, Xuan-Phi Nguyen, Austin Xu, Do Xuan Long, Minzhi Li, Chengwei Qin, Peifeng Wang, Silvio Savarese, and 1 others. 2025. *A survey of frontiers in llm reasoning: Inference scaling, learning to reason, and agentic systems*. *arXiv preprint arXiv:2504.09037*.

Zahra Khalila, Arbi Haza Nasution, Winda Monika, Aytug Onan, Yohei Murakami, Yasir Bin Ismail Radi, and Noor Mohammad Osmani. 2025. Investigating retrieval-augmented generation in quranic studies: A study of 13 open-source large language models. *arXiv preprint arXiv:2503.16581*.

Omer Nacar, Anis Koubaa, Serry Sibae, Yasser Al-Habashi, Adel Ammar, and Wadii Boulila. 2025. Gate: General arabic text embedding for enhanced semantic textual similarity with matryoshka representation learning and hybrid loss training. *arXiv preprint arXiv:2505.24581*.

Ossama Obeid, Nasser Zalmout, Salam Khalifa, Dima Taji, Mai Oudah, Bashar Alhafni, Go Inoue, Fadhl Eryani, Alexander Erdmann, and Nizar Habash. 2020. CAMEL tools: An open source python toolkit for Arabic natural language processing. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 7022–7032, Marseille, France. European Language Resources Association.

Islam Oshallah, Mohamed Basem, Ali Hamdi, and Ammar Mohammed. 2025. Cross-language approach for quranic qa. *arXiv preprint arXiv:2501.17449*.

Shuofei Qiao, Yixin Ou, Ningyu Zhang, Xiang Chen, Yunzhi Yao, Shumin Deng, Chuanqi Tan, Fei Huang, and Huajun Chen. 2023. Reasoning with language model prompting: A survey. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5368–5393, Toronto, Canada. Association for Computational Linguistics.

Stephen Robertson, Hugo Zaragoza, and 1 others. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.

Sander Schulhoff, Michael Ilie, Nishant Balepur, Konstantine Kahadze, Amanda Liu, Chenglei Si, Yinheng Li, Aayush Gupta, HyoJung Han, Sevien Schulhoff, and 1 others. 2024. The prompt report: a systematic survey of prompt engineering techniques. *arXiv preprint arXiv:2406.06608*.

Fanar Team, Ummar Abbas, Mohammad Shahmeer Ahmad, Firoj Alam, Enes Altinisik, Ehsannedin Asgari, Yazan Boshmaf, Sabri Boughorbel, Sanjay Chawla, Shammur Chowdhury, and 1 others. 2025. Fanar: An arabic-centric multimodal generative ai platform. *arXiv preprint arXiv:2501.13944*.

Cosine similarity; Top candidates = 200; Batch size = 8 for embedding computation

Cross-Encoder Reranking:

Model: Arabic BERT-based reranker; Final candidates = 5; Temperature = 0.1 for stable rankings

A.2 Prompt

You are an expert Islamic scholar. Your task is to answer multiple-choice questions. [Examples...] First, use the following reference text to determine the answer: RAG CONTEXT QUESTION: MULTIPLE CHOICES: Your response MUST be only the single capital letter of the correct option. Do not include 'Answer:', explanations, or any other text.

Chunking Strategy:

Target chunk size: 200 tokens (BM25 optimized); Overlap: 20 tokens between adjacent chunks; Minimum chunk size: 50 tokens; Maximum chunk size: 400 tokens

A Implementation - Technical Details

A.1 Hyperparameters

BM25 Configuration:

$k_1 = 1.2$ (term frequency saturation); $b = 0.75$ (length normalization); Top candidates = 1000

Dense Retrieval:

Embedding dimension = 768; Similarity metric =