

QIAS 2026: Overview of the Shared Task on Islamic Inheritance Reasoning

Abdessalam BOUCHEKIF¹, Somaya ELTANBOULY¹, Samer RASHWANI¹,
Shahd GABEN¹, Mutaz AL-KHATIB¹, Heba SBAHI¹,
Emad MOHAMED², Mohammed GHALY¹

¹Hamad bin Khalifa University, Qatar ²Nazarbayev University, Kazakhstan

{aboucchekif, seltanbouly, srashwani, sgaben, malkhatib, hsbahi, mghaly}@hbku.edu.qa
emad.mohamed@nu.edu.kz

Abstract

This paper presents a comprehensive overview of the QIAS 2026 shared task, organized as part of the OSACT7 Workshop and co-located with LREC 2026. The shared task was designed to evaluate the ability of large language models to perform complex reasoning in the religious and legal domain of Islamic inheritance. Unlike conventional question-answering benchmarks, QIAS 2026 focuses on end-to-end reasoning from natural language cases, requiring systems to perform the full inheritance calculation process, from identifying the eligible heirs to assigning the correct share to each beneficiary. To support this evaluation, the task was based on the MAWARITH benchmark, a dataset of 12,500 Arabic inheritance cases annotated with intermediate reasoning steps and final answers. System submissions were evaluated using MIR-E, a multi-step metric that measures performance across the main stages of inheritance reasoning. A total of 16 teams participated in the shared task, investigating a range of approaches, including prompting-based methods, retrieval-augmented generation, and fine-tuning strategies. The results show that Islamic inheritance remains a highly challenging benchmark for current language models, especially in stages that require precise legal interpretation and structured numerical reasoning. This overview summarizes the task design, dataset, evaluation framework, participating systems, and main results.

Keywords: multi-step reasoning, Islamic inheritance reasoning, Arabic language processing

1. Introduction

Large language models (LLMs) have recently achieved strong performance across a wide range of natural language processing tasks, including question answering, summarization, and complex text generation. Their success has been particularly evident in tasks that benefit from broad linguistic coverage and large-scale pretraining. However, LLMs still face important challenges in specialized domains that require precise reasoning, structured decision-making, and strict adherence to domain-specific rules. These limitations become more apparent in tasks that require a sequence of dependent reasoning steps, where an error at an early stage can propagate and compromise the final answer. This issue is especially important in religious and legal domains, where reasoning is not only knowledge-intensive but also constrained by formal principles and interpretive traditions.

In Islamic studies, and particularly in Islamic law, systems must reason over authoritative and highly structured sources such as the Qur'an, Hadith, and juristic writings. They must also operate within a framework based on clear legal principles and, in some cases, different interpretations across schools of thought. As a result, evaluating LLMs in such contexts requires benchmarks that go beyond surface-level question answering and instead test the ability of models to produce precise and

faithful reasoning. Islamic inheritance law (*ʿilm al-mawārith*) serves as an effective testbed for evaluating reasoning capabilities. This domain of Islamic jurisprudence requires multi-step legal and numerical reasoning to resolve cases. A valid solution must identify eligible heirs, determine which heirs are excluded or blocked, assign appropriate shares, assess the need for adjustments, and calculate the final distribution. The process is governed by strict jurisprudential rules, and certain cases present additional complexities such as *ʿawl* and *radd*. Due to its combination of legal interpretation, structured reasoning, and precise calculation, Islamic inheritance law offers a highly informative benchmark for assessing the reasoning abilities of modern language models. To support research in this direction, recent work introduced MAWARITH (Boucchekif et al., 2026), a large-scale benchmark of 12,500 Arabic inheritance cases designed for end-to-end reasoning from natural language descriptions. The benchmark includes detailed intermediate reasoning steps and final answers, making it possible to evaluate not only the correctness of the final output but also the validity of the reasoning process itself. This resource created an opportunity to move beyond standard multiple-choice evaluation and toward a more realistic setting in which systems must solve inheritance cases as humans would read and understand them.

The QIAS 2026 Shared Task was designed to eval-

uate whether participating systems can perform end-to-end Islamic inheritance reasoning in Arabic. It focuses on solving full cases from natural language, covering the complete reasoning process from heir identification to final share allocation. It also examines whether recent reasoning-oriented LLMs, such as Gemini, GPT, DeepSeek, Fanar, and Qwen, can transfer their strong performance on mathematical and synthetic benchmarks to the more complex domain of structured legal and religious reasoning.

In this paper, we present an overview of the QIAS 2026 Shared Task. We describe the task, the MAWARITH benchmark used in the evaluation, the MIR-E multi-step evaluation metric, the participating systems, and the main results obtained by the submitted approaches. We also discuss key lessons learned from the shared task and highlight the main challenges that current systems still face in Islamic inheritance reasoning.

2. Task Description

The QIAS 2026 Shared Task focuses on the end-to-end automation of Islamic inheritance reasoning (*‘ilm al-mawārith*) from natural language. The task requires systems to process Arabic inheritance cases and generate a complete, structured solution. For each input question, systems must produce a detailed step-by-step reasoning trace (`<think>`), followed by a concise final answer (`<answer>`). The task is formulated as a sequence of dependent reasoning stages. Systems are expected to identify and explicitly report the following components:

1. Identification of all mentioned heirs, including determining which are eligible or excluded based on applicable blocking rules (*hājib*).
2. Assignment of the correct legal shares (*furūd*) to entitled heirs according to classical Islamic inheritance jurisprudence, following the majority opinion (*al-jumhūr*).
3. Determination of whether a global adjustment to the distribution is required.
4. Computation of the final estate distribution, including cases where adjustments such as *‘awl* (proportional reduction) or *radd* (return of the residue) apply.

Participants must submit outputs in a structured format (e.g., JSON) that captures these reasoning steps. This representation enables fine-grained evaluation and allows analysis of different error types, such as legal reasoning errors versus numerical computation errors.

3. Data

The QIAS 2026 data come from the MAWARITH benchmark introduced by Boucekif et al. (2026).

It contains 12,500 cases written in Arabic and follows the majority opinion in Islamic inheritance law (*al-jumhūr*). Table 1 shows the distribution of cases across the training and test splits, as well as the complexity categories covered in the benchmark. Each case describes a complete inheritance situa-

Split	Simple	‘Awl	Radd	Total
Training	11,079	577	344	12,000
Test	456	39	5	500
Total	11,535	616	349	12,500

Table 1: Distribution of inheritance cases by legal complexity

tion in natural language. The system must identify the heirs mentioned in the case, determine which heirs are excluded by blocking rules, assign the correct legal shares, decide whether an adjustment is needed, and compute the final distribution of the estate. In this way, the benchmark does not test only final answers, but the full reasoning process required to solve an inheritance case correctly. The dataset covers a wide range of family relations found in classical Islamic inheritance law, including parents, children, spouses, siblings, grandparents, uncles, nephews, and other extended relatives. The MAWARITH dataset was built in several stages. First, inheritance cases were generated using the Almawarith calculator, which allows users to define heirs and their numbers through a structured interface and then produces the corresponding shares. This step provided a reliable foundation and helped ensure the correctness of the legal and numerical outcomes. Since the goal of MAWARITH is to evaluate reasoning from natural language, these structured cases were then rewritten as fluent Arabic inheritance questions, closer to real user queries. The outputs were then reviewed and enriched by an expert in Islamic studies, who added detailed legal and numerical explanations for each case. These explanations cover the main stages of inheritance reasoning, including heir identification, blocking rules, share assignment, and adjustment cases such as *‘awl* and *radd* when needed. To improve consistency, the expert-written explanations were standardized with the support of Gemini-2.5-Flash, while keeping the legal reasoning unchanged. Finally, the dataset was carefully validated to ensure consistency between the question, the reasoning steps, and the final inheritance shares.

3.1. MIR-E: Mawarith Inheritance Reasoning Evaluation

The QIAS 2026 shared task uses MIR-E (Mawarith Inheritance Reasoning Evaluation) (Boucekif et al., 2026), a weighted multi-stage metric for evaluating both intermediate reasoning steps and final

Team Name	Affiliations
CVPD (Swaileh et al., 2026)	CVPD Research group
Simplicity (Almansour, 2026)	King Saud University
KMS (Alkhamis, 2026)	King Saud University
QU-NLP (Alsmadi, 2026)	Qatar University
PSL (Mouhoub and Boucekif, 2026)	Paris Dauphine University
Silah (Kurdi et al., 2026)	Umm Al-Qura University
AGS-KSU (Sidaoui, 2026)	King Saud University

Table 2: Affiliations of teams that participated in the test phase and submitted a paper to QIAS 2026.

outputs in Islamic inheritance problems. Unlike standard evaluation based only on the final answer, MIR-E provides a more fine-grained assessment by scoring the main stages of the reasoning process. It includes four components:

1. **Heirs and Blocking** (S_h): evaluates whether the model correctly identifies the effective heirs, the blocked heirs, and their counts.
2. **Share Assignment** (S_s): measures whether the assigned shares for the eligible heirs are correct.
3. **Adjustment** (S_a): checks whether the model predicts the correct adjustment type (*none*, *awl*, or *radd*), and is scored only if the previous two stages are fully correct.
4. **Final Distribution** (S_f): evaluates whether the model produces the correct final distribution after the full inheritance process.

4. Results and Discussion

A total of 16 teams participated in the final phase. Table 2 summarizes the affiliations of the teams that participated in the test phase and submitted a paper. We provide a baseline implementation using Fanar-Sadiq, a modern Arabic large language model accessible via API. This baseline relies exclusively on prompting techniques, without any fine-tuning. The goal is to provide a simple yet effective reference point for evaluating model performance. The dataset and baseline code are publicly available online.¹ Overall, the submitted systems illustrate three main methodological directions: (i) end-to-end reasoning with prompting alone, (ii) fine-tuned models specialized for the task, and (iii) hybrid pipelines that combine LLM-based language understanding with deterministic symbolic reasoning. This diversity of approaches makes the shared task especially useful for comparing different strategies for Arabic legal reasoning. The participating systems explored different ways of solving Islamic inheritance cases. The most common approach explored

by participants was prompting-based end-to-end reasoning with large language models. Team PSL (Mouhoub and Boucekif, 2026) followed this setting and evaluated several models, including Gemini 2.5 Flash, Qwen3-32B, GPT-oss-120B, Llama-3.3-70B, Fanar-Sadiq, and Fanar-C-2-27B. Team KMS (Alkhamis, 2026) explored a similar direction and also evaluated Gemini 2.5 Pro and Mistral. Overall, the two teams reported similar performance trends: commercial models were generally more reliable, while open-weight models showed weaker results on this task.

Team CVPD (Swaileh et al., 2026), which achieved the best result in the shared task, proposed a RAG-based pipeline designed to generate outputs that match the MIR-E evaluation format. Their system built a knowledge base from synthetic legal question-answer pairs, with the option to include books, then retrieved relevant context for each Arabic inheritance question before generating a single JSON output. This output followed the required schema and included eligible heirs, blocked heirs, legal shares, the adjustment type when applicable, and the final *taṣīl* distribution. The system also included parsing and validation steps to ensure compatibility with the official evaluator. This approach showed that combining retrieval, controlled generation, and structured output constraints can be highly effective for this task.

Team Silah (Kurdi et al., 2026) investigated three strategies: retrieval-augmented generation based on a curated rule base, supervised fine-tuning of large language models, and a combination of fine-tuning and retrieval. Their experiments showed that fine-tuning alone performed better than retrieval-based approaches, with the best results obtained by a fine-tuned Fanar model. This suggests that task-specific fine-tuning can be an effective approach for Islamic inheritance reasoning. Team QU-NLP (Alsmadi, 2026) proposed a multi-stage Quantized Low-Rank Adaptation (QLoRA) strategy. Their approach involved initial domain adaptation on a corpus of Islamic fatwas to capture jurisprudential reasoning patterns, followed by task-specific fine-tuning on structured inheritance cases

¹https://gitlab.com/islamgpt1/qias_shared_task_2026

Rank	Team	Approach	MIR-E
1	CVPD	Hybrid RAG	0.935
2	Simplicity	Two-stage neuro-symbolic pipeline	0.931
3	KMS	Prompting-based end-to-end reasoning	0.916
4	QU-NLP	Multi-stage QLoRA fine-tuning	0.907
5	PSL	prompting-based end-to-end reasoning	0.898
6	AGS-KSU	QLoRA fine-tuning and Prompting-based end-to-end reasoning	0.84
7	Silah	LoRA fine-tuning	0.826
8	UTLM	Prompting-based end-to-end reasoning	0.742

Table 3: Official leaderboard results for QIAS 2026 for teams that participated in the test phase and submitted a paper

to optimize JSON-formatted output. This methodology allowed a relatively small 4B parameter model to achieve competitive performance, highlighting the effectiveness of specialized training strategies for complex legal reasoning tasks. Team AGS-KSU (Sidaoui, 2026) also explored a fine-tuning approach using Qwen2.5-3B, similar to the strategy adopted by Team QU-NLP. However, this fine-tuned model achieved a relatively low MIR-E score of 0.30. By comparison, their prompting-based configuration using GPT-5.4 Thinking reached 0.84, yielding a substantially stronger result.

Finally, some teams explicitly separated natural-language understanding from legal computation. Team Simplicity (Almansour, 2026) proposed a two-stage neuro-symbolic pipeline in which a commercial LLM was used only for Arabic information extraction. The extracted heirs were mapped to a standardized set of legal heir categories and then passed to a symbolic rule-based component that carried out blocking, share assignment, and final calculation according to the rules of *‘ilm al-mawāriṭh*. This design reflects a clear division of labor between the LLM and the symbolic module. The final performance of the submitted systems is summarized in Table 3, where Team CVPD achieved the top ranking (0.935) using a RAG pipeline based on the Qwen-9B model. Notably, the QU-NLP system’s multi-stage QLoRA approach achieved a competitive MIR-E score of 0.907. This performance is particularly significant, as it closely matches the 0.901 score of the commercial Gemini-2.5-Flash model (as reported in (Bouchekif et al., 2026)). These results demonstrate that, while RAG remains a powerful tool for precision, specialized domain adaptation can enable smaller open-weight models, such as the 4B-parameter Qwen3, to achieve performance levels close to those of commercial models on this task.

5. Conclusions and Future Work

In this paper, we presented the QIAS 2026 Shared Task, designed to evaluate the ability of large lan-

guage models to solve Islamic inheritance cases end-to-end, including case understanding, heir identification, rule application, and final share distribution. The results of the participating systems revealed a clear gap between commercial and open-weight models, with commercial systems generally achieving stronger performance on this challenging reasoning task. At the same time, the findings suggest that fine-tuning LLMs on MAWARITH data can substantially improve performance, while RAG also proved beneficial for some systems.

As a future direction, we plan to organize a follow-up edition of the shared task focusing on more complex inheritance cases, including pregnancy-related cases, multiple deaths, the missing person (mafqūd), and the intersex heir (khunthā). We also plan to focus on fine-tuning approaches for small Arabic LLMs. To ensure a fairer comparison across participants, future editions may restrict the use of commercial LLMs and focus on open-weight models that can be reproduced and compared under the same conditions.

6. Limitations

This work has several limitations. First, QIAS 2026 focuses only on inheritance cases based on the majority opinion (*al-jumhūr*). As a result, it does not represent the full diversity of juristic opinions in Islamic law. Second, although the benchmark covers many cases, it still does not include all the complexity of real inheritance situations. Some difficult or rare cases remain outside the current scope. Third, the MIR-E metric evaluates structured reasoning outputs in a detailed and reproducible way, but it does not fully measure explanation quality, clarity, or usefulness for real users.

Finally, some submitted systems use commercial models, which makes full reproducibility more difficult. Future work should expand the legal coverage of the benchmark and give more attention to open and reproducible systems.

7. References

- Ummar Abbas, Mohammad Shahmeer Ahmad, Minhaj Ahmad, Abdulaziz Al-Homaid, Anas Al-Nuaimi, Enes Altinisik, Ehsaneddin Asgari, Sanjay Chawla, Shammur Chowdhury, Fahim Dalvi, et al. 2026. Fanar 2.0: Arabic generative AI stack. *arXiv preprint arXiv:2603.16397*.
- Ali Abdelaal, Mohammed Nader Al Haffar, Mahmoud Fawzi, and Walid Magdy. 2026. [Islam-icmmlu: A benchmark for evaluating llms on islamic knowledge](#). *arXiv preprint 2603.23750*.
- Nouar AIDahoul and Yasir Zaki. 2025. NYUAD at QIAS shared task: Benchmarking the legal reasoning of LLMs in arabic islamic inheritance cases. In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 861–866.
- Omar Alkhamis. 2026. KMS at QIAS 2026: Evaluating llm reasoning for islamic inheritance division. In *Proceedings of the 7th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT7), co-located with LREC-COLING 2026*, Palma de Mallorca, Spain.
- Muath Almansour. 2026. Simplicity at QIAS 2026: Decoupling language extraction from mathematical logic in islamic inheritance law. In *Proceedings of the 7th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT7), co-located with LREC-COLING 2026*, Palma de Mallorca, Spain.
- Ameera Almasoud, Sharefah Al-Ghamdi, Reem Alqifari, Noof Alfear, and Hend Al-Khalifa. 2026. [MirathQA: A dataset for evaluating large language models on hanbali islamic inheritance reasoning tasks](#). *Data in Brief*, page 112589.
- Sanaa Alowaidi. 2025. [SEA-Team at QIAS 2025: Enhancing LLMs for question answering in islamic texts](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 940–946, Suzhou, China. Association for Computational Linguistics.
- Mohammad Alsmadi. 2026. QU-NLP at QIAS 2026: Multi-stage QLoRA fine-tuning for arabic islamic inheritance reasoning. In *Proceedings of the 7th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT7), co-located with LREC-COLING 2026*, Palma de Mallorca, Spain.
- Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Salah Eddine Bekhouche, Abdellah Zakaria Selam, Hichem Telli, Cosimo Distanto, and Abdenour Hadid. 2025. [CVPD at QIAS 2025 shared task: An efficient encoder-based approach for islamic inheritance reasoning](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 929–934.
- Gagan Bhatia, Hamdy Mubarak, Mustafa Jarrar, George Mikros, Fadi Zaraket, Mahmoud Alhirthani, Mutaz Al-Khatib, Logan Cochrane, Kareem Darwish, Rashid Yahiaoui, et al. 2026. From RAG to agentic RAG for faithful islamic question answering. *arXiv preprint arXiv:2601.07528*.
- Abdessalam Boucekif, Shahd Gaben, Samer Rashwani, Somaya Eltanbouly, Mutaz Al-Khatib, Heba Sbahi, Mohammed Ghaly, and Emad Mohamed. 2026. [MAWARITH: A dataset and benchmark for legal inheritance reasoning with LLMs](#). *arXiv preprint arXiv:2603.07539*.
- Abdessalam Boucekif, Samer Rashwani, Emad Soliman Ali Mohamed, Mutaz Alkhatib, Heba Sbahi, Shahd Gaben, Wajdi Zaghouni, Aiman Erbad, and Mohammed Ghaly. 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: Shared Tasks*, pages 851–860.
- 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*, pages 246–257.
- Iheb Chaabane, Puneesh Khanna, Suhail Mohamad, Slim Frikha, Shi Hu, Abdalgader Abubaker, Reda Alami, Mikhail Lubinets, Mohamed El Amine Seddik, Hakim Hacid, et al. 2026. Falcon-H1R: Pushing the reasoning frontiers with a hybrid model for efficient test-time scaling. *arXiv preprint arXiv:2601.02346*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *CoRR*, abs/2110.14168.
- DeepSeek AI. 2024. DeepSeek-R1: Incentivizing reasoning capability in large language models.

- Technical report, DeepSeek AI. Technical Report.
- Eman Elrefai, Mohamed Lotfy Elrefai, and Aml Hassan Esmail. 2025. [Gumball at QIAS 2025: Arabic LLM automated reasoning in islamic inheritance](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 953–959, Suzhou, China. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#).
- Shehenaz Hossain and Haithem Afli. 2025. [ADAPT–MTU HAI at QIAS2025: Dual-expert LLM fine-tuning and constrained decoding for arabic islamic inheritance reasoning](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 923–928, Suzhou, China. Association for Computational Linguistics.
- Ghader Kurdi, Hanan Justanieah, and Hala Justanieah. 2026. [Silah at QIAS 2026: Fine-tuning vs. retrieval-augmented generation for islamic inheritance reasoning](#). In *Proceedings of the 7th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT7), co-located with LREC-COLING 2026*, Palma de Mallorca, Spain.
- AL-Smadi Mohammad. 2025. [QU-NLP at QIAS 2025 shared task: A two-phase llm fine-tuning and retrieval-augmented generation approach for islamic inheritance reasoning](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 892–898.
- Mohamed Motasim Hamed, Nada Ghneim, and Riad Sonbol. 2025. [HIAST at QIAS 2025: Retrieval-augmented LLMs with top-hit web evidence for arabic islamic reasoning QA](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 883–891, Suzhou, China.
- Mohamed Lamine Mouhoub and Chahinez Boucekif. 2026. [PSL at QIAS 2026: Which models perform better in arabic inheritance reasoning?](#) In *Proceedings of the 7th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT7), co-located with LREC-COLING 2026*, Palma de Mallorca, Spain.
- Hamdy Mubarak, Rana Malhas, Watheq Mansour, Abubakr Mohamed, Mahmoud Fawzi, Majd Hawasly, Tamer Elsayed, Kareem Mohamed Darwish, and Walid Magdy. 2025. [IslamicEval 2025: The first shared task of capturing LLMs hallucination in islamic content](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 480–493.
- Yossra Noureldien, Hassan Suliman, Farah Attallah, Abdelrazig Mohamed, and Sara Abdalla. 2025. [Athar at QIAS2025: LLM-based question answering systems for islamic inheritance and classical islamic knowledge](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 914–922, Suzhou, China. Association for Computational Linguistics.
- Nguyen Xuan Phuc and Thìn Đăng Văn. 2025. [PuxAI at QIAS 2025: Multi-agent retrieval-augmented generation for islamic inheritance and knowledge reasoning](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 905–913.
- Jihad R'baiti, Chouaib El Hachimi, Youssef Hmamouche, and Amal El Fallah Seghrouchni. 2025. [MorAI at QIAS 2025: Collaborative LLM via voting and retrieval-augmented generation for solving complex inheritance problems](#). In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 947–952, Suzhou, China. Association for Computational Linguistics.
- Weizhou Shen, Ziyi Yang, Chenliang Li, Zhiyuan Lu, Miao Peng, Huashan Sun, Yingcheng Shi, Shengyi Liao, Shaopeng Lai, Bo Zhang, et al. 2025. [QwenLong-L1.5: Post-training recipe for long-context reasoning and memory management](#).
- Hicham Ghazi Sidaoui. 2026. [AGS-KSU at QIAS 2026: A comparative study of prompting and LLM approaches for structured islamic inheritance reasoning](#). In *Proceedings of the 7th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT7), co-located with LREC-COLING 2026*, Palma de Mallorca, Spain.
- Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, et al. 2025. [OpenAI GPT-5 system card](#). *arXiv preprint arXiv:2601.03267*.
- Wassim Swaileh, Mohammed-En-Nadhir Zighem, Hichem Telli, Salah Eddine Bekhouche, Abdellah Zakaria Sellam, and Fadi Dornaika. 2026. [CVPD at QIAS 2026: RAG-guided LLM reasoning for al-mawarith share calculation and heir allocation](#). In *Proceedings of the 7th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT7), co-located with LREC-COLING 2026*, Palma de Mallorca, Spain.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Jesse Woo, Fateme Hashemi Chaleshtori, Ana Marasović, and Kenneth Marino. 2025. BriefMe: A legal NLP benchmark for assisting with legal briefs. *arXiv preprint arXiv:2506.06619*.

Osama Farouk Zaki. 2025. CIS-RG at QIAS 2025 shared task: Approaches for enhancing performance of LLM on islamic legal reasoning and its mathematical calculations. In *Proceedings of The Third Arabic Natural Language Processing Conference: Shared Tasks*, pages 935–939.

A. Related Work

Large language models have recently been applied to a wide range of Islamic knowledge tasks, including Quranic question answering (Bhatia et al., 2026), knowledge retrieval (Mubarak et al., 2025; Phuc and Vãn, 2025), and the analysis of hallucinations in Islamic content (Mubarak et al., 2025). A recent related benchmark is *IslamicMMLU* (Abdelal et al., 2026), which evaluates LLMs on broad Islamic knowledge across the Quran, Hadith, and Fiqh through a large multiple-choice benchmark. These studies show that LLMs perform well on knowledge retrieval and basic understanding when answers rely on direct textual matching. However, they often hallucinate and show clear limitations on tasks that require structured reasoning or deep domain knowledge. In (Boucekif et al., 2025b), the authors report that several models, such as LLaMA and ALLaM, frequently cite non-existent Quranic verses or fabricate Hadith references. As a result, the generated conclusions are not only incorrect, but are also supported by false religious evidence. This behavior raises serious concerns for religious and legal applications, where correctness depends not only on the final answer but also on the authenticity and reliability of the cited sources.

To mitigate hallucinations, recent work has explored Retrieval-Augmented Generation (RAG) approaches. While RAG improves access to relevant information and enhances factual faithfulness and citation accuracy (Bhatia et al., 2026; Noureldien et al., 2025; Alowaidi, 2025), it remains insufficient for answering questions that require multi-step reasoning. This limitation has motivated the development of reasoning-oriented models that explicitly support multi-step inference. In this context, recent research has increasingly focused on such

models that aim to move beyond surface-level text generation and promote more reliable reasoning.

Models such as *o3*, *GPT-5* (Singh et al., 2025), *Gemini-2.5* (Anil et al., 2023), *Gemini3*, *DeepSeek-R1* (DeepSeek AI, 2024), along with open models such as *Fanar-C-2-27B* (Abbas et al., 2026), *Falcon-H1R* (Chaabane et al., 2026), *Fanar-Sadiq* (Abbas et al., 2026) and *Qwen3* (Shen et al., 2025) illustrate this trend by promoting more consistent multi-step inference through instruction tuning and large-scale pretraining. Evaluations of reasoning-oriented language models have largely focused on mathematical and logical benchmarks, on which these models have achieved strong results, particularly in arithmetic reasoning, symbolic manipulation, and competition-style mathematics (Cobbe et al., 2021; Hendrycks et al., 2021; Wei et al., 2022). Beyond mathematical and logical benchmarks, recent work has begun to investigate LLM reasoning in legally grounded settings by evaluating models on legal benchmarks such as BRIEFME (Woo et al., 2025), which require structured argumentation and rule-based reasoning. The authors show that *GPT-4o* can outperform human annotators on argument summarization by producing clear and coherent summaries. Even within the Islamic domain, inheritance law has received growing attention as a challenging testbed for LLM reasoning (AIDahoul and Zaki, 2025; R’baiti et al., 2025; Zaki, 2025). In particular, QIAS 2025² (Boucekif et al., 2025a) was a shared task dedicated to Islamic inheritance law (*‘ilm al-mawārith*), focusing on the evaluation of large language models under strict, rule-based legal and numerical constraints, using a benchmark of 2,200 MCQs. A similar MCQs benchmark is *MirathQA* (Almasoud et al., 2026), built from 1394 inheritance cases. Studies report that commercially deployed (e.g., *Gemini* and *ChatGPT*), reasoning-oriented models consistently outperform non-reasoning or general-purpose models on benchmarks requiring multi-step inference and structured reasoning (Boucekif et al., 2025b; Mohammad, 2025; Bekhouche et al., 2025; Motasim Hamed et al., 2025; Hossain and Afli, 2025). Additionally, Elrefai et al. (2025) show that a fine-tuned Qwen3 model achieved top-ranked performance on the QIAS 2025 shared task. However, this evaluation setup does not allow assessing whether models truly reason correctly. Models were required to select a single correct answer among six options, without any evaluation of the validity of their intermediate reasoning steps or the correctness of the legal justifications leading to that choice. Moreover, Boucekif et al. (2025b) shows that even when a model selects the correct answer, the underlying reasoning can still be incorrect or legally invalid. In contrast, the present shared task

²<https://sites.google.com/view/qias2025/>

requires models to perform end-to-end inheritance reasoning, explicitly generating intermediate reasoning steps, applying jurisprudential rules, and computing the final inheritance shares.