



The Assembly of Muslim Jurists of America
22nd Annual Imams' Conference
Dallas – United States

Can AI Be a Mufti or Judge?

Legal Formalism and Juristic Insight in Islamic Law

Dr. Yousef Wahb

Yousef Wahb is the Director of Scholar Development and Academic Partnerships at Yaqeen Institute, a Research Fellow in Philanthropy and Islamic Law at Huron University, and an author for LexisNexis Canada, specializing in Family Law and Islamic Finance. He holds an LL.M. from Windsor Law, an M.A. from the University of Chicago Divinity School, and a B.A. in Islamic Studies from Al-Azhar University. He is currently pursuing a Ph.D. in Islamic Law in the Middle Eastern Studies program at the University of Chicago

"الآراء في هذا البحث تعبر عن رأي الباحث وليس بالضرورة عن رأي أئمة"

Opinions in this research are solely those of the author and do not represent AMJA.




TABLE OF CONTENTS

1. Introduction	4
1.1 Distinguishing Iftā' from Qaḍā'	5
1.2 Scope	6
2. Framework: Formalism, Realism, and the Islamic Parallel	8
2.1 The Spamann-Klöhn-Posner Experiments	8
2.2 The Paradox of the Official Story	10
2.3 From Non-Mechanical Discernment to Fiqh al-Nafs	11
3. The Conditions of the Muftī and the Judge	13
3.1 Meaning, Application, and the Limits of Retrieval	13
3.2 Fiqh al-Nafs	15
3.3 The Mufti as Trustee: Al-Qarāfī's Reminders to Muftis	18
4. 'Illa, Qiyās, and the Limits of Algorithmic Reasoning	21
4.1 'Illa Identification as Evaluative Judgment	21
4.2 The Epistemic Problem: AI Cannot Verify Its Own Reasoning	23
Conclusion	24

"I fell ill with a sickness in which I forgot everything I had memorized, even the Qur'an! But I did not forget *fiqh*, because my knowledge of everything besides *fiqh* was knowledge of memorization, whereas my knowledge of *fiqh* was knowledge of guidance, like a man who was absent from his city for a time, then returned: would he forget the way to his own house?"¹

— The Abbasid Caliphate Chief Judge Abū Yūsuf al-Anṣārī (d. 182/798)

"I predict that human judges will be around for a while."²

— Chief Justice John G. Roberts, Jr.

1. Introduction

These two statements, spoken twelve centuries and two legal traditions apart, converge on the same insight. Law is not merely a body of information to be stored and retrieved. It is a faculty, a way of seeing and reasoning, cultivated within the human person and cannot be reduced to what is documented or encoded in a dataset.

Abū Yūsuf lost his memorized knowledge, the transmitted material, but not his legal knowledge. To him, *fiqh* was an internalized faculty so deeply woven into his being that illness could not dislodge it. It was, in his own metaphor, like knowing the way home. When Chief Justice Roberts was asked whether AI might one day replace human judges, his answer rested on a similar intuition. Something essential to the judicial function resists mechanization. The classical scholars had a name for what Abū Yūsuf retained: juristic insight (*fiqh al-nafs*). Al-Juwaynī (d. 478/1085) called it "the capital of the independent jurist (*mujtahid*)," adding that it "cannot be attained merely by memorizing books."³ Modern AI systems, whether language models that generate text or classifiers that map input into predefined categories, learn statistical patterns in large datasets, retaining and reproducing those patterns rather than any faculty like the one Abū Yūsuf described. It is neither understanding nor guidance. They have no such faculty, no way home.

Generative AI can now produce outputs that closely resemble fatwas and judicial reasoning. Large language models (LLMs) trained on vast corpora of text can generate responses to *fiqh* questions that are fluent, well-organized, and superficially plausible. But resemblance raises a foundational question: *can there be a fatwa without a mufti, or a hukm without a judge?* The question is not primarily technological, since AI will almost certainly improve. The central issue is jurisprudential: What does Islamic law require of those who exercise its legal functions? The conditions of valid *iftā'* and adjudication belong to the mufti and the judge, and not the fatwa or the ruling by itself. If an entity fails to meet those conditions, then even a perfectly worded fatwa is, in the classical framework, not a fatwa at all.

This paper argues that generative AI cannot independently perform the functions of fatwa-issuing (*iftā'*) or Islamic adjudication (*qadā'*), and does so on three grounds. First, it draws on recent empirical research in American legal theory by Holger Spamann, Lars Klöhn, and Eric Posner, which tested how federal judges, law students, and AI decide the

¹ Reported in Muḥammad Zāhid al-Kawtharī, *Ḥusn al-taqāḍī fī sirat Abī Yūsuf al-Qāḍī* (Maktaba al-Azhariyya li-l-Turāth, ___), 14–17.

² Roberts, J. G. (2023). *2023 year-end report on the federal judiciary*. Supreme Court of the United States. <https://www.supremecourt.gov/publicinfo/year-end/2023year-endreport.pdf>

³ 'Abd al-Malik b. 'Abdullāh b. Yūsuf al-Juwaynī, *al-Burhān fī uṣūl al-fiqh*, ed. 'Abd al-'Aẓīm al-Dīb, 2 vols. (Qatar, 1978), 2:1332–1333.

same case. The experiment showed that AI is rigidly formalist. That formalism speaks to a distinction Islamic legal theory cares about as well, between mechanical rule-application and cultivated juristic judgment. AI performs certain tasks within legal reasoning well, such as collecting precedents, reviewing documents, and amalgamating primary and secondary legal sources, but that competence does not extend to the evaluative judgment this paper is concerned with.

Second, it examines the classical conditions of the mufti and the judge. AI fails two of these: the jurisprudential condition of the cultivated human faculty of *fiqh al-nafs* and the accountability condition that presupposes responsibility before God and the community. Neither condition can be reduced to data or computation. That jurisprudential condition is developed further below through the semantics/pragmatics distinction in the philosophy of language and through Gadamer's account of interpretive application, both of which show why training on text, however extensive, does not close the gap between knowing what words mean and knowing what a specific case or speaker requires. Third, it tests current AI models against the technical operations of legal reasoning, specifically the identification of the operative legal cause (*'illa*) and the performance of analogical reasoning (*qiyās*). Even on this narrower, more "mechanical" ground, AI falls short.

The paper analyzes classical definitions of "jurist" (*faqīh*) and "legal reasoning" (*ijtihād*), examines the qualifications of muftis and judges, and tests the classical condition of juristic intuition and skill (*malaka*) against AI. An algorithm is data-driven and it does not think; it mimics. It generates statistically probable sequences of words without access to meaning, moral agency, or accountability. Training a model on Islamic legal texts does not make it a jurist any more than feeding a machine the writings of experienced judges makes it a judge.⁴

1.1 Distinguishing Iftā' from Qaḍā'

This paper addresses AI's role in two distinct legal functions that carry different conditions, consequences, and implications for AI's involvement. *Iftā'* is the issuance of a non-binding legal opinion (fatwa) by a qualified scholar (mufti) in response to a question. It is advisory in nature, addressed to the questioner's conscience, and does not carry the force of law. The mufti is described as one who signs on behalf of God. While questioners are free to act or seek another opinion, the mufti's reporting of God's command establishes grounds for religious accountability. Legal theorists characterize the fatwa as a "report" (*ikhbār*) in the sense that the mufti translates God's universal law into a specific context for the

⁴ The philosophical debate over whether AI systems can "understand" is extensive, from Searle's Chinese Room argument (1980) to more recent critique such as Bender et al.'s "stochastic parrots" (2021) and Dennett's "competence without comprehension" (2017). The P-Zombie problem poses the question most directly. An entity that produces outputs indistinguishable from those of a conscious being, but with no inner experience, is what al-Qarāfī would call a document without a *ḥukm* behind it. For an exploration of this question from an Islamic philosophical perspective, see Ramon Harvey, "Can AI Be Conscious? An Islamic Philosophical Perspective," *Oases of Wisdom*, June 5, 2026. This paper does not attempt to resolve these debates. Its argument rests on Islamic legal theory's own criteria, which presuppose qualities the philosophical literature broadly agrees AI systems currently lack.

questioner.⁵ But even translation requires intimate familiarity with the universal principles and particular rules of exception, the options of diverging or converging with established precedents or practices, the ability to detect hidden motives in a questioner's framing, and sensitivity to local custom.

Qaḍā' is binding adjudication: the issuance of a legally enforceable ruling (*ḥukm*) by a judge in a dispute. It resolves conflicts, establishes rights, enforces obligations, and carries institutional consequences. Unlike the fatwa, which reports the law, the judicial ruling is an "origination" (*inshā'*): the judge brings a new, specific obligation into existence for the litigants. This origination requires delegated authority from God through the political sovereign. That status presupposes moral agency and accountability, conditions rooted in personhood.⁶ An algorithm cannot originate a binding moral reality because origination, in al-Qarāfī's (d. 684/1285) framework, is a volitional one, not a textual one. The judicial ruling is "mental speech" (*kalām nafsānī*), an internal act that subsists within the mind of the judge. The written document or spoken pronouncement is a record of that internal act. The medium of communication (whether written, spoken, or gestured) does not affect the ruling's validity because the ruling itself is the judge's internal determination, not the external expression. An LLM operates entirely at the level of external expression. It produces text. Its processes, while computationally complex, are not the kind of internal act al-Qarāfī describes: a volitional determination by a conscious moral agent. The judge's volition here is not a separate condition alongside his reading of the case. It is what makes that reading his, an act for which he, and no one else, answers. The same is true of the mufti's *fiqh al-nafs*, which is the mode in which interpretation becomes accountable to God, not a technical skill added on top of it. The stakes and institutional consequences of *iftā'* and *qaḍā'* differ, but the underlying requirement does not. In both, interpreting the law and answering for that interpretation are a single act, not two. This is the paper's central claim that AI, lacking a volitional and accountable self behind its output, cannot supply that single, unified act.

1.2 Scope

The scope of this paper is limited to AI in the seat of authority, that is, AI performing or purporting to perform the end-to-end functions of the mufti and the judge. The question is not whether AI can perform any given step within *iftā'* or *qaḍā'* in isolation; later sections identify specific steps, such as retrieving texts and precedent, where AI performs well. The question is whether AI can independently drive the process as a whole, from the questioner's or litigant's initial presentation to the justified, binding ruling, without a human jurist supplying the judgment at each stage. The broader questions of AI as a research tool for scholars, AI in Islamic education, and AI in community applications fall outside this paper's focus.

AI is already embedded in legal and religious practice. A 2026 Northwestern University study found that over 60% of US federal judges use at least one AI tool in their work, primarily for legal research and document review, though fewer than 5% use it to

⁵ Shihāb al-Dīn Aḥmad b. Idrīs al-Qarāfī, *al-Iḥkām fī tamyīz al-fatāwā 'an al-aḥkām wa-taṣarrufāt al-qāḍī wa-al-imām*, ed. 'Abd al-Fattāḥ Abū Ghudda (Cairo: Dār al-Salām, 2009), xx; English translation in *The Criterion for Distinguishing Legal Opinions from Judicial Rulings and the Administrative Acts of Judges and Rulers*, trans. Mohammad H. Fadel (New Haven, CT: Yale University Press, 2017).

⁶ Al-Qarāfī, *al-Iḥkām*, __

inform decisions.⁷ Islamic institutions are similarly engaged. The UAE's Islamic Affairs department and Egypt's Dar al-Ifta have introduced automated systems to handle religious queries, and Jordan's Ifta Department has cautioned against reliance on AI-produced fatwas.⁸ AI's apparent competence in legal reasoning has also drawn direct empirical study. A recent blinded evaluation found that law professors preferred LLM-generated answers to contracts-law questions over their own peers' answers in a large majority of trials.⁹ Results of this kind are sometimes read as evidence that AI has already closed the gap this paper is concerned with. Section 4 below addresses why that reading does not hold up once the object of measurement is examined closely. This paper does not argue against the use of AI as a tool. It argues against the conflation of tool-use with the exercise of legal authority.

⁷ "Majority of US federal judges are using AI, study finds," *Reuters* (2026) <https://www.reuters.com/legal/government/majority-us-federal-judges-are-using-ai-study-finds-2026-03-30/>

⁸ See Sophia Tsourlaki, "Artificial Intelligence and Fatwā Issuance: A Case Study of Dubai and Egypt," *Islamic Inquiries* 1, no. 2 (2022): 107–125. Tsourlaki concludes that IACAD's system simply identified lexical proximity between user questions and pre-registered fatwas, functioning more as a search tool than as artificial intelligence in any substantive sense. It does not gain significant adoption among Muslims. See also 'Amr Sayyid 'Imām "al-Dhakā' al-iṣṭinā'ī wa-ṣinā'at al-fatwā: qirā'a fī al-ḍawābiṭ wa-l-āthār," Paper presented at Dar al-Ifta Conference, Cairo, 2025. <https://conference.fatwaacademy.org/Content/data/themes/02/02-18.pdf>; Muḥammad al-Za'ātreh, "Athar al-dhakā' al-iṣṭinā'ī fī ṣiyāghat al-fatwā," *Dār al-Iftā' al-Urdunniyya*, June 18, 2023. <https://www.aliftaa.jo/article/5683/>.

⁹ Salinas et al., "Law Professors Prefer AI Over Peer Answers" (Stanford University et al., May 2026). See the discussion in §4 below of what this kind of preference-based evaluation does and does not measure.

2. Framework: Formalism, Realism, and the Islamic Parallel

American jurisprudence has been shaped for over a century by a tension between two accounts of what judges actually do.¹⁰ Legal formalism holds that correct outcomes can be derived from rules and precedent through logical application. The judge's role is mechanical: identify the relevant rule, apply it to the facts, and reach the outcome the law dictates. Chief Justice Roberts captured this at his 2005 confirmation hearing when he compared judges to baseball umpires who "call balls and strikes" but "don't make the rules."¹¹

Legal realism, the dominant counter-tradition since the early twentieth century, holds that where rules underdetermine the outcome, the gap is filled by factors realism locates outside the formal legal materials: a judge's experience, moral intuition, and sense of the case's practical stakes. Rules are often ambiguous, and facts rarely map neatly onto legal categories. On a more careful view, the written opinion is not simply a rationalization of a prior, unconstrained decision. It reflects a disciplined but extra-textual judgment, what Karl Llewellyn called "situation sense." Through that judgment, professional training and experience narrow the range of legally acceptable outcomes even where the rules alone do not.¹² This debate matters for AI because the difference between formalist and realist decision-making becomes visible in these hard cases, where formal rules run out. A large language model absorbs whatever methodology is present in its training data, including realist reasoning, but cannot perform the act of contextual judgment that realism requires. It can mimic the language of contextual sensitivity without exercising it. This is also why Spamann, Klöhn, and Posner built their experiment around a genuinely contested war crimes appeal rather than a routine case. The formalism/realism distinction, and AI's placement on one side of it, is only detectable at this margin.

2.1 The Spamann-Klöhn-Posner Experiments

¹⁰ I use "formalism" and "realism" here as descriptive labels for a behavioral spectrum: rigid rule-adherence versus context-sensitive departure from precedent. This is not a claim about the historical positions of any jurisprudential school. The popular version of this debate, which casts formalism as the denial of judicial discretion, is a caricature drawn largely from realism's own polemics; see Brian Z. Tamanaha, *Beyond the Formalist-Realist Divide: The Role of Politics in Judging* (Princeton, NJ: Princeton University Press, 2010). The genuine disagreement concerns the degree and legitimate sources of interpretive discretion. See Frederick Schauer, "Formalism," *Yale Law Journal* 97, no. 4 (1988): 509–548; Brian Leiter, "Positivism, Formalism, Realism," *Columbia Law Review* 99, no. 4 (1999): 1138–1164. Posner and Saran's own use of "formalistic" to describe GPT-4o's behavior tracks this same descriptive sense: a pattern of rigid adherence to precedent and insensitivity to case-specific factors. It is not an endorsement of one jurisprudential school as historically or theoretically correct.

¹¹ U.S. Senate, Committee on the Judiciary, *Confirmation Hearing on the Nomination of John G. Roberts, Jr. to Be Chief Justice of the United States*, 109th Cong., 1st sess., September 12, 2005.

¹² Karl N. Llewellyn, *The Common Law Tradition: Deciding Appeals* (Boston: Little, Brown, 1960). Contrast the earlier, more theatrical version of this claim in Joseph C. Hutcheson Jr., "The Judgment Intuitive: The Function of the 'Hunch' in Judicial Decision," *Cornell Law Quarterly* 14 (1929): 274–288.

The empirical evidence strongly favors the realist account and reveals that AI falls squarely on the formalist side. The experimental lineage begins with Holger Spamann and Lars Klöhn (2016), who designed an experiment based on a real case before the International Criminal Tribunal for the Former Yugoslavia. Thirty-two US federal judges were asked to decide a simulated appeal involving a military commander charged with aiding and abetting war crimes. The experiment used a 2×2 factorial design (meaning two variables were manipulated simultaneously, each with two levels) producing four conditions. The first variable was the degree of sympathy toward the defendant. One version presented a sympathetic Croatian who expressed remorse and worked toward reconciliation after the war, while the other presented an unsympathetic Serb who ridiculed the tribunal and showed no remorse. The second variable was the direction of precedent: a prior ruling favoring conviction, or a prior ruling favoring acquittal. Each judge was randomly assigned to one of the four conditions.

The results were striking. Judges were strongly influenced by sympathy, the legally irrelevant variable. Affirmations rates differed by 45% between the sympathetic and unsympathetic defendants. But they were not significantly influenced by precedent. When they wrote their opinions, not a single judge mentioned sympathy. Every opinion cited precedent, statutory interpretation, and standard legal reasoning. The factors that actually drove their decisions were absent from their written justifications. Spamann and Klöhn called this the gap between the “official story” of judging and its reality.¹³

In a follow-up study, Spamann and Klöhn (2024) repeated the experiment with 91 law students at three top-ranked US law schools. The results were the opposite. Students followed precedent rigidly and were unmoved by sympathy. In other words, students were formalists while judges were realists.¹⁴

Eric Posner and Shivam Saran (2026) then replicated the experiment using GPT-4o, the most advanced commercially available large language model at the time. They ran each of the four conditions approximately 50 times to account for the randomness in how LLMs generate text. GPT-4o behaved like the students, not like the judges. It followed precedent closely and was unaffected by sympathy. It was, in Posner’s words, “formalistic.”¹⁵

Three aspects of Posner’s findings are especially relevant. First, prompt engineering failed to bridge the gap. Posner and Saran tried multiple strategies to coax GPT-4o into behaving more like human judges, instructing it to act as a “legal realist,” telling it to consider the defendant’s sympathetic qualities, and framing the task as a social science experiment rather than a judicial one. None produced a consistent sympathy effect. As Posner concludes, “something about how LLMs are designed and operate will frustrate prompt-engineering techniques that attempt to spur LLMs to decide cases as judges really decide them rather than according to the public reasons that judges provide in their

¹³ Holger Spamann and Lars Klöhn, “Justice Is Less Blind, and Less Legalistic, Than We Thought: Evidence from an Experiment with Real Judges,” *Journal of Legal Studies* 45, no. 2 (2016): 255–80.

¹⁴ Holger Spamann and Lars Klöhn, “Can Law Students Replace Judges in Experiments of Judicial Decision-Making?,” *Journal of Law and Empirical Analysis* 1, no. 1 (2024): 149–161.

¹⁵ Eric A. Posner and Shivam Saran, “Judge AI,” *Journal of Legal Analysis* 16, no. 1 (2026): 235–262.

opinions.”¹⁶ Second, a robustness test across numerous other LLMs largely confirmed the GPT-4o results.¹⁷ The formalist tendency is not specific to one model or one company; it held across every model Posner tested.¹⁸ Third, in a subsequent lecture on the findings, Posner posed a pointed question: “Would you want law students to be judges?”¹⁹ Both students and LLMs are competent at applying rules and can produce coherent legal reasoning. Neither exercises the contextual judgment that distinguishes a seasoned judge from a student reading a casebook.

The relevance of these findings to Islamic legal theory is not that Muslim jurists were “legal realists” in the modern American sense. The experiments demonstrate a distinction essential to both traditions between mechanical rule-application and cultivated juristic judgment.

2.2 The Paradox of the Official Story

Posner’s most pointed observation is what he calls the paradox of the official story: judges are influenced by factors they do not acknowledge in their written opinions.²⁰ The written opinion is the official account of why a case was decided. It may be incomplete, or even inaccurate account, about what actually drove the decision. The gap between official and actual legal reasoning is widely recognized in legal theory, though American jurisprudence itself is divided over its significance. H. L. A. Hart’s positivism treats officials’ outward convergence in practice as sufficient for a rule of recognition to exist. Their private motivations for following do not matter.²¹ Even Baude and Sachs, who defend the official story’s authority on Hartian grounds, acknowledge that officials may “pay lip service” to shared legal accounts while privately adhering to different reasoning.²² This is not the only position within American legal theory. Ronald Dworkin’s account of law as integrity denies that lip service is compatible with valid adjudication. A judge whose decision is not actually controlled by the legal materials has failed at judging, regardless of how the opinion is written.²³ The Hart/Dworkin divide shows that whether genuine internal commitment is required, or merely convergent behavior, is contested even within the tradition that

¹⁶ Posner tested off-the-shelf GPT-4o, not a model fine-tuned for judicial reasoning. A fine-tuned model might produce outcomes that more closely resemble those of human judges, but Posner and Saran argue that any such improvement would more plausibly reflect pattern-matching than genuine contextual judgment. Posner and Saran, “Judge AI,” 7.

¹⁷ Some newer models showed a weak sympathy effect in some configurations, but the effect was inconsistent, occasionally reversed in direction, and in every case dwarfed by the precedent effect. Posner and Saran, “Judge AI,” 17.

¹⁸ Posner and Saran’s follow-up study confirmed this pattern with newer models on a different judicial experiment. See Eric A. Posner and Shivam Saran, “Silicon Formalism,”

¹⁹ Eric A. Posner, “Judge AI” (Nora and Edward Ryerson Lecture, University of Chicago, April 16, 2026). Reported at <https://www.law.uchicago.edu/news/eric-posner-explores-promise-and-limits-ai-judging-ryerson-lecture>.

²⁰ Posner and Saran, “Judge AI,” 17-9.

²¹ H. L. A. Hart, *The Concept of Law*, 3rd ed. (Oxford: Oxford University Press, 2012), ch. 6, on the “internal point of view” and the rule of recognition. Hart’s account allows that officials may accept a rule of recognition for widely differing reasons, including calculation of self-interest, so long as their practice converges.

²² William Baude and Stephen E. Sachs, “The Official Story of the Law,” *Oxford Journal of Legal Studies* 43, no. 1 (2023): 178-201.

²³ Ronald Dworkin, *Law’s Empire* (Cambridge, MA: Harvard University Press, 1986), esp. ch. 7 on “law as integrity.”

produced the official-story literature. What is distinctive about Islamic legal theory is narrower. It ties the legitimacy of interpretive variance to whether it traces to a cultivated faculty sincerely exercised on the sources. Conformity to a shared account is not enough, on its own, without inquiry into whether that account actually produced the ruling.

The published fatwa is the formal articulation of a ruling. It states the question, identifies the relevant evidence, and announces the ruling. But behind every fatwa is a process of verifying how the general ruling applies to this specific questioner in these specific circumstances (*tahqīq al-manāṭ*). That process involves the mufti's assessment of the questioner's condition, their compliance capacity, the customary practice of their community, and other contextual factors that rarely appear in the written fatwa itself. This variability is not, on the classical tradition's own account, evidence of a defect. Al-Ghazālī observes that the evidentiary indicators (*qarā'in*) bearing on a ruling are not self-interpreting, and that jurists reach different conclusions from the same evidence because their capacity to perceive what the evidence indicates differs, and even changes within the same jurist over time.²⁴ The tradition's complaint is not against variance as such, which it anticipates and, within limits, accepts as the natural consequence of *ijtihād* exercised by different, genuinely engaged jurists. Its complaint is against variance that does not trace to a sincerely exercised faculty at all, whether because the interpreter never engaged the sources or because his conclusion was driven by unaccountable inclination (*hawā*) rather than by the evidence.

One might argue that training on the full scholarly corpus, not just published rulings, gives AI access to the deliberative framework behind *iftā'*. But this confuses describing a methodology with executing it. This is the problem Posner identified. GPT-4o had access to the entire legal realist literature, and still decided cases like a formalist. Knowing the theory of judgment did not produce judgment. Sherman Jackson's account of *uṣūl al-fiqh* points to a related problem. The jurist approaches the sources with prior orientation, shaped by experience and context. The formal apparatus of *uṣūl* then tests and disciplines against the evidence.²⁵ That formal apparatus is where AI's training data is concentrated. It has no orientation of its own to bring to that testing process, only patterns drawn from how others have written about theirs.

2.3 From Non-Mechanical Discernment to Fiqh al-Nafs

In the American legal context, the finding that judges are influenced by sympathy is treated as significant primarily because it goes undisclosed. The sentencing and equity traditions already permit judges considerable discretion to weigh a defendant's circumstances. But the Spamann experiment shows sympathy operating on appellate judgment while remaining entirely absent from the written opinion's stated grounds.²⁶ That gap is what generates the debate, not the mere presence of sympathy as a factor.

²⁴

²⁵ Sherman A. Jackson, "Fiction and Formalism: Toward a Functional Analysis of Uṣūl al-Fiqh," in *Studies in Islamic Legal Theory*, ed. Bernard G. Weiss (Leiden: Brill, 2002), 177–201.

²⁶ The debate is also inseparable from a specifically American anxiety about judicial legitimacy. If judges decide on undisclosed grounds, they arguably exercise an authority the separation of powers does not grant them (the countermajoritarian difficulty). Islamic legal theory, which treats *ijtihād* as a licensed and obligatory juristic function rather than a usurpation of legislative power, does not share this particular source of unease, even as it independently requires that a ruling be traceable to its actual grounds.

The realist response is that sympathy is inevitable and perhaps even beneficial, because rigid rule-application can produce unjust outcomes. But neither side treats sympathy as a formal requirement of the judicial role. The claim is not that the judges' susceptibility to sympathy in the Spamann experiment is desirable, nor that Islamic law endorses it. Islamic law explicitly condemns judging by personal inclination (*hawā*). A judge who acquits a defendant because of ethnic sympathy would be derelict in his duty. The experiment reveals that competent adjudication involves a non-mechanical element of discernment that formalist accounts cannot capture. The question is how that discernment is treated.

In Islamic legal theory, this broader capacity for contextual sensitivity, which sympathy imperfectly expresses in the American case, is not treated as an accident of human nature. It is a methodological requirement known as *fiqh al-naḥs*, inseparable from the obligation to verify that a general ruling actually applies to the particular case at hand. Jackson argues that *uṣūl al-fiqh* reaches its limits at the point of contextual indicators (*qarā'in*), at least to the extent it operates as a formal system of interpretive rules.²⁷ Even on his reading, *uṣūl* provides no instruction on how hard, how long, or where to look for contextual indicators that might qualify a general command. The level of scrutiny a jurist applies is driven by concerns external to the formal rules of interpretation. Whatever one makes of Jackson's broader thesis about the relationship between *uṣūl* and legal reasoning, this particular observation holds. The search for context is one the formal rules alone do not resolve.²⁸ That space is where the jurist's trained judgment operates, and a machine restricted to formal rules has no way to navigate it.

AI's inability to exercise *fiqh al-naḥs* is not an incidental limitation, nor a feature that might be engineered away with better models or more data. Even if a model were fine-tuned to generate outputs similar to a jurist exercising *fiqh al-naḥs*, the outputs would only reflect "learned statistical associations between the case condition and the judges' decisions" rather than genuine contextual judgment.²⁹

²⁷ Jackson, "Fiction and Formalism," 190-8??

²⁸ While I draw on Jackson's analysis to highlight where the formal rules of *uṣūl* reach their limits, I do not adopt his broader characterization of legal theory as primarily rhetorical. I subscribe to the view that Islamic legal theory is structurally non-formalist: the vast majority of rulings are probable rather than certain, and the system presupposes ongoing interpretive effort at every level. If the law were purely mechanical, there would have been no need for jurists to theorize at all.

²⁹ Posner and Saran, "Judge AI," 7.

3. The Conditions of the Muftī and the Judge

The previous section established that Islamic legal theory, like the legal realist tradition, treats contextual judgment as essential to the legal function, not as an incidental byproduct of human nature. This section develops the theoretical basis for that claim, and turns to the specific conditions that classical scholars placed on those who exercise *iftā'* and *qaḍā'*.

Islamic law places its legal functions in the hands of the *mukallaf*, the morally and legally responsible person. *Taklīf* requires three conditions: rational capacity, maturity, and volition. The ensuing legal capacity (*ahliyya*) is of two types: the capacity to bear rights and obligations (*ahliyyat al-wujūb*) and the capacity to perform legal acts (*ahliyyat al-adā'*). The classical scholars broadly required a core set of conditions of the mufti, which typically included Islam, *taklīf*, scholarly competence including *fiqh al-naḥs*, probity (*adāla*), and being from the people of *ijtihād*. The precise enumeration varies across the schools, and certain questions, such as the validity of a fatwa issued by a *fāsiq* or the status of the mufti *muqallid*, remain debated. The conditions of the judge are at least as stringent. Even in schools that relaxed some of these conditions, moral agency and accountability were never among them. AI is not the kind of subject these conditions describe. It cannot be addressed by revelation, cannot intend an act, and cannot answer for how it arrived at an output, because there is no one there to answer. The failure is structural rather than technical: the classical conditions are conditions of personhood, not performance.

3.1 Meaning, Application, and the Limits of Retrieval

The claim that AI lacks discernment, moral perception, or lived social understanding invites an obvious objection. These sound like present limitations, the kind that improve with better models, more data, or careful fine-tuning. To show that the limitation is principled rather than technical, this section explains what these capacities actually are.

Linguists distinguish sentence meaning from speaker meaning. Sentence meaning is the conventional, context-independent content a sentence carries by virtue of its words and grammar. Speaker meaning is what a particular speaker, in a particular context, intends to communicate by using those words. This content routinely diverges from, exceeds, or narrows the sentence's literal content.³⁰ A large language model is built to model the first. Its training optimizes for the statistical regularities of linguistic form: given this sequence of words, what continuation is probable. It has no access to the second, because speaker meaning is not a property of a text at all. It is a property of a mind, recovered by another mind through inference about intention, background, and purpose. Recovering it requires an interpreter who can model *this* speaker, not language in general.³¹

Al-Qarāfī's account of *fiqh al-naḥs*, discussed below, is best understood as an extended meditation on exactly this gap. But the point can be seen already in a much older illustration. Shaykh al-Azhar al-Sharqāwī (d. 1227/1812) describes al-Shāfi'ī instructing his student al-Muzanī: "*Witr* prayer is not obligatory, based on what has been established in my view; and intention in prayer is obligatory, based on what has been established in my view." Al-Muzanī can reproduce these rulings and argue for them, according to his teacher's

³⁰ H. P. Grice, "Logic and Conversation," in *Syntax and Semantics 3: Speech Acts*, ed. Peter Cole and Jerry L. Morgan (New York: Academic Press, 1975), 41–58.

³¹ Dan Sperber and Deirdre Wilson, *Relevance: Communication and Cognition*, 2nd ed. (Oxford: Blackwell, 1995).

doctrine, at the level of general categories: the claim-establishing factor (*muqtaḍī*) and the counter-indicator (*nāfi*).³² But as Ibn Abī Sharīf (d. 906/1501) observes, this is merely general evidence (*dalīl ijmālī*): al-Muzanī has the label, not the specific evidence that constitutes it, and cannot defend the ruling when challenged to say what the *muqtaḍī* or *nāfi* actually is in this case.³³ The gap between the two is not a gap in information. Al-Muzanī has the same sentences al-Shāfiʿī has. What he lacks is the situated competence, developed through prolonged engagement with a real interlocutor, to specify how those sentences apply and to hold that application against a challenge. A large language model sits at exactly al-Muzanī's level in this analogy, and for a structurally similar reason. It can reproduce rulings, label supporting principles, and present the positions of different schools with considerable fluency, because all of this is sentence-level content. What it cannot do is what al-Muzanī cannot do: supply and defend the specific, case-anchored evidence that only emerges through engagement with a particular challenge from a particular interlocutor. Training a model on the entire corpus of legal theory does not close this gap, for the same reason training GPT-4o on the entire legal realist literature did not produce a realist judge. More text does not convert sentence-level competence into speaker-level, situated competence.

This has a further consequence, developed influentially by Hans-Georg Gadamer. Understanding a legal or textual norm cannot be separated from applying it to a concrete situation. That application is achieved dialogically, through an interpreter's real engagement with the parties and particulars before him, not through prior textual analysis alone.³⁴ Al-Qarāfi's own account of *iftā'*, examined in the sections below, illustrates this: the facts relevant to a ruling are frequently not given in advance. They are drawn out through the mufti's questioning of the petitioner, a real exchange between two people, one of whom is exercising discernment about the other in real time. The relevant "text" a ruling responds to does not yet exist when the question is first asked. It is co-produced through the encounter. The classical tradition makes a version of this same point about the jurist himself, independent of any particular case. A jurist is properly called a *faqīh* only once he is disposed toward *fiqh* (*al-mutahayyi'*), and disposition is not yet the thing itself. Even a scholar who has mastered *uṣūl* and internalized the sources is not, in the tradition's own terms, a *faqīh* until he actually produces legal knowledge by engaging the evidence.³⁵ Readiness is not exercise. This is why training a model on the complete corpus of Islamic

³² 'Abdullāh b. Ḥijāzī al-Sharqāwī, *Ḥāshiyat al-Sharqāwī 'alā sharḥ al-taḥrīr wa ma'ahā taqrīrāt al-'allāma al-ṭallāwī al-musammā Fath al-qadīr*, 12 vols. (Amman: Dār al-Fath, 2025), 1:338-9.

³³ Al-Sharqāwī, *Ḥāshiyat*, 339. Al-Ghazālī took this further. He argued that the rationales of legal rulings were understood by the early authorities without recourse to dialectical methods. He warned that excessive engagement with disputation weakens juridical intuition (*dhawq al-fiqh*), as sound legal judgment often depends on forms of discernment that transcend formal argumentation. Abū Ḥāmid al-Ghazālī, *Iḥyā' ulūm al-dīn*, 10 vols. (Jedda: Dār al-Minhāj, 2011), 1:238. If disputation, itself a structured human activity, can erode legal discernment, the implications are clear for AI, which operates entirely at the level of formal pattern-matching.

³⁴ Hans-Georg Gadamer, *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), Part II.II.2, "The Hermeneutic Significance of Application."

³⁵ See Jalāl al-Dīn al-Maḥallī, *al-Badr al-ṭālī 'bi sharḥ Jam' al-jawāmi'*, ed. Murtaḍā al-Dāghistānī, 2 vols. (Beirut: Risala Foundation, 2012), 2:379.

legal literature does not make it a jurist. A fully trained model is, at best, analogous to a fully prepared scholar who has not yet exercised the faculty, except that a model has no faculty to exercise in the first place; it can only ever occupy the position of readiness, never the position of the act.

This also answers an objection that might otherwise seem to limit the argument's reach. One might grant that AI cannot perform *ijtihād*, the derivation of a novel ruling. But one could still suggest that AI functions adequately at the level of *taqlīd*, applying an already-settled rule to a case, since this is what most *iftā'* actually involves and where most questioners seek answers. But applying a settled rule to a case is not a mechanical act of subsumption. It requires determining that the facts before you actually satisfy the rule's conditions. Al-Qarāfī's Cairo case, discussed below, demonstrates this operation: a paradigmatically settled question ("is a marriage contract concluded in Cairo valid?"), whose correct disposition still depended entirely on relational, dialogical discernment of the petitioner's real circumstances and purpose. If even *taqlīd*-level *iftā'* requires this, the argument of this section applies to the whole of *iftā'*, not only to its *ijtihād*-level instances.

3.2 Fiqh al-Nafs

Among the conditions of the jurist (be it a mufti or a judge), *fiqh al-nafs* stands apart. It is, in al-Juwaynī's formulation, the "guiding framework" that governs the exercise of all the others:

Sound juristic insight (*fiqh al-nafs*) is required, for it is the capital of the *mujtahid*. It cannot truly be acquired through effort alone; if one is naturally endowed with it, that is what is intended, otherwise it cannot be attained merely by memorizing books... Knowledge of *uṣūl* is indispensable, sound juristic insight is the guiding framework.³⁶

Al-Maḥallī (d. 864/1459) defines a jurist who has *fiqh al-nafs* as one possessing "a deep, natural understanding of what legal language intends" (*shadīd al-fahm bi-l-ṭab' li-maqāṣid al-kalām*), adding that "without it, the derivation intended by *ijtihād* cannot be achieved."³⁷ Al-ʿAṭṭār (d. 1250/1834) clarifies that this faculty is rooted in natural disposition rather than formal instruction. It is directed specifically at legal language, excluding, for example, Sufi-style allusions (*ishārāt*), which do not qualify as *fiqh*.³⁸

Al-Juwaynī describes *fiqh al-nafs* as partly innate: a quality refined through training and experience, but whose foundation is a God-given capacity.³⁹ Taqī al-Dīn al-Subkī (d. 756/1355) adds that the *mujtahid* is one for whom the legal sciences have become a settled faculty (*malaka*), who has "encompassed the majority of the rules of the Sharia and practiced them such that he acquired a capacity by which he understands the intent of the Lawgiver."⁴⁰ The *mujtahid*'s competence comes from prolonged practice over time. *Malaka* is developmental, experiential, and irreducibly human. The later tradition, visible in al-Maḥallī, al-ʿAṭṭār, and al-Subkī, shows an impulse to specify further criteria for what this faculty consists in. It registers an anxiety about intuition's resistance to definition that al-

³⁶ Al-Juwaynī, *al-Burhān*, 2:1332–1333.

³⁷ Al-Maḥallī, *al-Badr*, 2:380.

³⁸ Ḥasan al-ʿAṭṭār, *Ḥāshiya ʿalā Jamʿ al-jawāmiʿ*, 2 vols. (Cairo: Dār al-Baṣāʾir, 2009), 2:422. Al-ʿAṭṭār derives this from the verb *faqīha*, which belongs to the category of *afʿāl al-sajāyā* (verbs denoting innate character traits).

³⁹ Al-Juwaynī, *al-Burhān*, 2:1332–1333

⁴⁰ Al-Maḥallī, *al-Badr*, 2:381.

Juwaynī and al-Ghazālī did not share. Al-Juwaynī simply asserts that *fiqh al-naḥs* is the *mujtahid*'s "capital" and cannot be attained by memorization alone, without feeling obliged to further pin down what the faculty technically consists in. The scholastic effort to define it more precisely is a later development.

Fiqh al-naḥs is intuition. It is a way of perceiving what legal language intends, and what a case actually requires, that cannot be reduced to a rule or a checklist. But this intuition is a trained one. Continued engagement with the sources, and not raw impression, is what refines it. And it is trained on legal language specifically, with sentiment alone unable to produce it. This is also what *ijtihād* itself is. It is an art, exercised by a trained person, not a science carried out by a procedure that gives the same output for the same input every time. An algorithm is exactly that kind of procedure: fixed steps applied uniformly to whatever text it receives. *Fiqh al-naḥs*, on every description above, is the opposite kind of thing. If it cannot be acquired from books alone, as al-Juwaynī insists, it cannot be acquired from a model trained on books either.

Two bodies of modern scholarship, one philosophical and one historical, help bring the shape of this kind of faculty into sharper view. Neither offers a new checklist for identifying *fiqh al-naḥs*; that would repeat the very anxiety just described. Both instead explain why a faculty like this takes the form it does: acquired through practice, resistant to codification, and irreducible to a stored set of rules.

Pierre Bourdieu's concept of *habitus* describes a similar structure in a modern idiom. *Habitus*, for Bourdieu, is an acquired system of dispositions, produced through immersion in a structured field. It generates practice through the orientations and limits it places on a trained capacity for improvisation; the process is anything but the "mechanical determinism" Bourdieu explicitly contrasts it with.⁴¹ Applied specifically to law, Bourdieu develops the notion of a juridical *habitus*: the practical mastery that trained jurists acquire through professional formation. This mastery lets them generate legally appropriate reasoning without consciously deriving it from first principles on each occasion.⁴² This is neither memorized rule-following nor unconstrained intuition. It is what al-Subkī describes: a capacity produced by encompassing and practicing the rules until they become a settled disposition through which the jurist perceives what a case requires. The borrowing here is narrow. This paper draws on *habitus* as a description of how a non-mechanical, trained legal faculty is formed and exercised, not on Bourdieu's broader sociological project of treating legal reasoning as reducible to competition for position and authority within a juridical field.⁴³

The historical-anthropological record confirms that traditional Islamic legal pedagogy was itself organized around producing exactly this kind of *habitus*, not around transmitting

⁴¹ Pierre Bourdieu, "Structures, Habitus, Practices," in *The Logic of Practice*, trans. Richard Nice (Stanford, CA: Stanford University Press, 1990), 52–65, esp. 95 on *habitus* as generating practice through "the mediation of the orientations and limits it assigns to the *habitus*'s operations of invention," rather than mechanical determinism.

⁴² Pierre Bourdieu, "The Force of Law: Toward a Sociology of the Juridical Field," trans. Richard Terdiman, *Hastings Law Journal* 38, no. 5 (1987): 805–853.

⁴³ Bourdieu's broader argument in "The Force of Law" is in some tension with this paper's premises. Part of his project is to show juridical practice as shaped by competition for symbolic capital and position within the juridical field, a dynamic that in the same readings comes at the expense of the field's claims to autonomous rational justification. This paper draws only on his account of *habitus* as a trained, non-mechanical generative disposition, without adopting that broader deflationary thesis.

propositional content. Aria Nakissa's ethnographic and textual study of religious education at al-Azhar documents that traditional legal training centered on apprenticeship with a shaykh, imitation of the teacher's reasoning and comportment, and internalization of mastery through practice rather than through the transfer of discrete information. Ibn Khaldūn (d. 808/1406) himself theorized this model explicitly in terms of *malaka*.⁴⁴ This paper draws on Nakissa's documentation of that pedagogical structure without adopting the further Foucauldian thesis, derived from Talal Asad, that legal knowledge is therefore reducible to the discursive and pedagogical apparatus that produces it. That thesis is difficult to square with *uṣūl al-fiqh*'s own self-understanding as a theory of evidence, answerable to the epistemic status of sources and to a ruling's traceability under challenge, rather than a self-contained practice whose only standard is its own internal reproduction. The *ribā* case study examined below is a direct counterexample. Shāfi'ī jurists revised their identification of the operative cause over centuries precisely because economic reality changed, not because the discursive tradition required renewal for its own sake.⁴⁵

Fiqh al-naḥs manifests in specific, practical capacities that the classical scholars described in detail. Ibn al-Qayyim (d. 751/1350) warns that whoever gives fatwas "merely based on what is transmitted in books, despite differences in their customs, habits, times, places, circumstances, and contextual indicators, has gone astray and led others astray. His harm to the religion is greater than the harm of a physician who treats all people, despite differences in their lands, habits, times, and temperaments, according to a single medical book."⁴⁶ Part of the mufti's competence is knowing when not to answer. Where AI systems have been engineered to refuse certain queries, these refusal behaviors reflect design choices by developers, not the moral restraint that al-Juwaynī describes. The mufti's decision to withhold a ruling is itself an exercise of *fiqh al-naḥs*; an AI system's refusal is a programmed constraint.

Al-Khaṭīb al-Baghdādī (d. 463/1071) elaborates on what contextual awareness requires: "The jurist needs to have some familiarity with every matter of this world and the Hereafter, and knowledge of seriousness and jest, agreement and opposition, benefit and harm, the affairs that commonly occur among people, and their recognized customs."⁴⁷

AI can access none of this in the way contemplated by the classical jurists, and the gap does not stem from present technical limitation. As the preceding discussion shows, its competence is at the level of sentence meaning and, at best, statistical pattern. It can detect linguistic patterns associated with tone or intention, because these are statistical

⁴⁴ Aria Nakissa, "An Epistemic Shift in Islamic Law: Educational Reform at al-Azhar and Dār al-'Ulūm," *Islamic Law and Society* 21, no. 3 (2014): 209–251, esp. 220–227 on habitus-based transmission of legal and linguistic mastery, drawing on Ibn Khaldūn's theorization of both as forms of *malaka* acquired through companionship (*ṣuḥba*) with a teacher.

⁴⁵ Nakissa's account is embedded in a broader Foucauldian framework, adopted from Talal Asad's notion of Islam as a "discursive tradition," which treats legal knowledge as constituted by the practices and pedagogical structures that produce it. This paper does not adopt that further thesis. *Uṣūl al-fiqh* has its own theory of evidentiary weight, its procedures for verifying that a conclusion is licensed by a text, and its requirement that a ruling be defensible under challenge. That apparatus presupposes a jurist answerable to something beyond the discursive practice itself.

⁴⁶ Ibn Qayyim al-Jawziyya, *I'lām al-muwaqqi'īn 'an rabb al-'ālamīn* (Beirut: Dār al-Jīl, n.d.), 3:78.

⁴⁷ Al-Khaṭīb al-Baghdādī, *al-Faqīh wa-l-mutafaqqih*, ed. 'Ādil ibn Yūsuf al-'Azzāzī (Riyadh: Dār Ibn al-Jawzī, 2000), __

regularities in text. *Fiqh al-nafs*, on every classical description surveyed here, operates at the level of speaker meaning and situated application, a cultivated intuition exercised on a specific person in a specific case, not a pattern extracted from text in general. It cannot bear responsibility for the consequences of a ruling, nor can it exercise prudential restraint grounded in moral accountability, since accountability presupposes a self capable of being called to account, and none is present here.

3.3 The Mufti as Trustee: Al-Qarāfī's Reminders to Muftis

Al-Qarāfī devoted the final section of *The Criterion* to ten practical "Reminders" for muftis, each addressing a dimension of the mufti's work that goes beyond textual knowledge.⁴⁸ Three of these are especially instructive for the question of AI, because together they show that *iftā'* is fundamentally relational and intersubjective. Each Reminder identifies a point where the correct ruling depends on an exchange between two people rather than on the correct processing of a text. *Fiqh al-nafs*, as Section 3.2 describes it, is a hermeneutical faculty aimed at legal language, perceiving what a text or an evidence actually indicates. What al-Qarāfī describes here is that same order of discernment, cultivated, non-mechanical, resistant to codification, turned toward a person instead of a text. The object changes; the kind of faculty required does not.

The First Reminder concerns the mufti's ability to distinguish between two kinds of intention behind a questioner's language: a confirming intention, where the speaker means what the words literally say, and a particularizing intention, where the speaker means to exclude part of what the words literally cover. Al-Qarāfī illustrates this with the case of someone who swears, "By God, I shall not wear any garment today." If the oath-taker later says he intended only linen garments, the mufti must determine whether this was a genuine intent to exclude other garments from the oath (a particularizing intention) or simply that only linen occurred to him at the time (a confirming intention that covers only part of the meaning). The legal consequences differ significantly. Al-Qarāfī warns that most muftis get this wrong. They ask "What were you thinking of at the time of the oath?" and take the answer at face value. The real question is different: did the oath-taker mean to conceptualize linen to the exclusion of everything else the words cover, or did linen simply happen to be what came to mind, with no intention to exclude anything else?⁴⁹ This is, in the terms introduced above, a pure question of speaker meaning: the sentence "any garment" is semantically stable, but what this particular speaker meant by it is not recoverable from the sentence at all.

This is a linguistic judgment that depends on the mufti's understanding of how ordinary people use language, what they are likely to have considered when speaking, and what their stated intentions actually reveal about their mental state. An AI system can process the sentence's semantic content, the range of things "any garment" could denote. But the mufti's task here is a question of pragmatics, recovering what this particular petitioner intends, not what the sentence could mean in general. An AI system has no such property to detect, in the sentence or anywhere else.

In the Fifth Reminder, al-Qarāfī states that the mufti "should not take the words of a member of the public at their face value when he asks for a legal opinion until he ascertains his intent." When a petitioner uses language or asks about a topic inconsistent with his personal background, "that constitutes a cause for suspicion which should cause the mufti

⁴⁸ Al-Qarāfī, *al-Iḥkām*, 227-68.

⁴⁹ Al-Qarāfī, *al-Iḥkām*, 227-9.

to investigate the petitioner's background to determine the petitioner's true identity and circumstances, and not to rely solely on the petitioner's formulation of the question."⁵⁰

Al-Qarāfī illustrates with a case from his own practice. He was asked about the validity of a marriage contract concluded in Cairo. The question seemed innocuous, but al-Qarāfī grew suspicious and refused to answer until the petitioner explained why he was asking, "for everyone knows that a marriage contract concluded in Cairo is valid." After persistent questioning, the petitioner revealed that the real purpose was to circumvent the legal prohibition against a woman remarrying her previous husband (*tahlīl* marriage). Al-Qarāfī ruled it impermissible regardless of location.⁵¹ Notice what kind of case this is. The stated question was not a hard case of *ijtihād*; the validity of Cairene marriage contracts was settled law, known to everyone. And yet the correct ruling depended entirely on facts that did not exist at the moment the question was asked. They were produced by al-Qarāfī's persistent questioning, a real dialogue with a real petitioner. There was no text to retrieve that already contained the *tahlīl* purpose; that purpose only entered the record because a human mufti pressed a human petitioner until it surfaced.

The mufti's competence lay less in knowing the answer to the stated question than in recognizing that the stated question was not the real question. That recognition required suspicion, investigation, and the kind of interpersonal discernment that al-Qarāfī calls a duty of the mufti's office. An AI-advocate might object that this is a problem of missing information; give the system enough background on the petitioner and it could flag the same suspicion. But this misses the point that the petitioner's actual purpose did not exist anywhere for a system to retrieve, background or otherwise, until al-Qarāfī's questioning drew it out. No amount of prior data closes a gap that only a real exchange between two people could close, because the fact in question was produced by that exchange, not hidden somewhere before it. Al-Qarāfī's competency lay in bearing responsibility for pursuing the suspicion and refusing to answer until the truth was disclosed.

The Seventh Reminder raises concerns about the mufti's awareness of pretexts. Al-Qarāfī warns that "many rights are sought for the sake of injustice." He gives the example of a usurer who asks, "Is it permissible to sell merchandise for cash?" If the mufti says yes, the usurer will use the ruling to justify selling a trifle and a silver coin in exchange for one thousand silver coins. The question is technically legitimate; the intent behind it is not. Al-Qarāfī advises that the mufti should either qualify his response to prevent the harmful consequences or refuse to answer entirely, and considers the latter the better course. He then invokes the prophetic tradition: "Jurists are God's trustees over His creation."⁵² The mufti's trusteeship is not confined to getting the legal answer right. It extends to that person's welfare, moral welfare included, and to whether a correct answer, honestly given, will end up doing harm. The mufti is not a dispenser of legal information. He is a trustee who exercises discretion about what to disclose, to whom, and when.

Taken together, these three Reminders describe a practice of *iftā'* that is irreducibly interpersonal. They share a single underlying structure. The mufti reads the questioner's language for unstated intentions (Reminder 1), investigates when the question seems suspicious (Reminder 5), and guards against pretexts that would use a correct answer for an unjust purpose (Reminder 7). Each is an expression of the same responsibility: a trusteeship over the petitioner as a person, not only over the correctness of a legal answer.

⁵⁰ Al-Qarāfī, *al-Iḥkām*, 236-7.

⁵¹ Al-Qarāfī, *al-Iḥkām*, 237.

⁵² Al-Qarāfī, *al-Iḥkām*, 241-2.

None of these operations appears in the published fatwa. All of them are, in al-Qarāfī's view, obligatory duties of the mufti's office. They are the same kind of cultivated, non-mechanical discernment *fiqh al-nafs* names for legal language, exercised here on persons instead of texts. And they constitute the gap between the "official story" of the fatwa and the reality of the process that produced it.

American corporate law is currently working through a parallel question, in a wholly secular register. Lawyers have begun ejecting AI note-taking bots from client meetings before privileged discussion begins, out of concern that the bot's presence, not any human's indiscretion, may itself void attorney-client privilege.⁵³ Courts are divided on the underlying premise. One federal judge ruled that a defendant's use of an AI system for legal advice could not be privileged, since the system's own privacy policy disclosed user data to third parties, while another declined to treat a litigant's use of a different AI tool as equivalent to disclosure to an outside party. Both rulings rest on the same premise that an AI system is not a person capable of standing inside a legal relationship. They disagree only about which of two distinct doctrines that premise controls, whether it forecloses privilege from ever attaching, or whether it forecloses treating disclosure to the tool as disclosure to an adversary. Applying a settled premise to a specific doctrinal question was itself a judgment call made by the human interpreter, the same structure this paper has argued *iftā'* and *qaḍā'* require. A confidential, trust-based legal relationship is not reducible to the information exchanged within it. Whether a non-human party's mere presence changes the relationship's legal character is a live question for American courts, for the same reason it is a live question for *iftā'*. In both cases, the relationship, not only its content, is what the law is protecting.

⁵³ Sarah Kessler, "All Those A.I. Note Takers? They're Making Lawyers Very Nervous," *The New York Times*, May 9, 2026. The article discusses conflicting rulings, one by Judge Jed S. Rakoff (S.D.N.Y.) in *United States v. Heppner* and one by Magistrate Judge Anthony P. Patti (E.D. Mich.) in *Warner v. Gilbarco, Inc.*, No. 2:24-cv-12333, 2026 WL 373043 (E.D. Mich. Feb. 10, 2026), both issued in February 2026, on whether use of an AI system voids attorney-client privilege or its functional equivalent.

4. 'Illa, Qiyās, and the Limits of Algorithmic Reasoning

The previous sections argued that AI cannot meet the conditions of the mufti or the judge, and that it lacks the cultivated interpersonal faculty of *fiqh al-naḥs*. This section addresses a different question: even if we set aside the conditions of the person, are there *technical operations* of Islamic legal reasoning that AI can perform? Some steps, such as retrieving relevant texts, identifying which scholars addressed a question, and presenting competing positions, are well within AI's capabilities. Others require evaluative judgment that cannot be reduced to retrieval or pattern-matching.

The structured account of legal reasoning in the Islamic tradition might seem to favor the mechanization argument. Al-Ghazālī attributes to al-Shāfi'ī a procedure for the *mujtahid* confronting a new case: first examine the explicit texts of the Qur'an, then the widely transmitted (*mutawātir*) hadith, then the solitary (*āḥād*) reports, then the apparent meanings of the Qur'an, then scholarly consensus, and only then resort to *qiyās*. This looks like a flowchart an algorithm could follow. But each step requires judgment that the procedure itself does not supply. Al-Shāfi'ī himself, in this account, concludes the sequence by warning that the *mujtahid* "does not resort to mechanical extension (*ṭard*) if he believes in God and knows the sources of the law."⁵⁴ The procedure is a framework for organizing judgment, not a substitute for it.

Legal analogical reasoning involves four elements: the original case (*aṣl*) for which a ruling exists in the texts, the new case (*far'*) for which no explicit ruling exists, the ruling (*ḥukm*) of the original case, and the operative legal cause (*'illa*) that links the two. The validity of the analogy depends entirely on the correct identification of the *'illa*, the most contested step in the entire structure of *qiyās*. Legal theorists developed elaborate systems for it: *takhrīj al-manāṭ* (extracting the operative cause), *tanqīḥ al-manāṭ* (refining it by eliminating irrelevant attributes), and *taḥqīq al-manāṭ* (verifying that it applies in the new case). Each of these steps involves judgment calls that the formal rules do not fully determine.

4.1 'Illa Identification as Evaluative Judgment

Consider what *'illa* identification actually requires. When the question of organ donation arose in the twentieth century, scholars had to determine whether donating a kidney is analogous to selling a body part (and therefore prohibited) or analogous to saving a life at personal cost (and therefore permitted under the principle of preserving life). The answer depends on which *'illa* one identifies for the original prohibition. If the *'illa* is commodification, donation may be permissible. If the *'illa* is bodily harm without necessity, even donation is questionable. If the *'illa* is violation of God's ownership over the body, both sale and donation are problematic. Similarly, when the Prophet ﷺ was asked about a man who had sexual intercourse during the day in Ramadan and ordered him to expiate,⁵⁵ the jurist must determine which features of the case are legally relevant. Is the ruling tied to the individual? To the act of intercourse? To Ramadan specifically? Or to the violation of

⁵⁴ Abū Ḥāmid al-Ghazālī, *al-Mankhūl min ta'liqāt al-uṣūl*, ed. Muḥammad Hīto (Damascus: Dār al-Fikr, 1998), 466-7.

⁵⁵ Ṣaḥīḥ al-Bukhārī, no. 1936.

fasting per se? Different answers produce different legal structures. The act of isolating the operative cause depends on judgments about relevance that the text does not determine.

An AI system would retrieve the various scholarly positions and organize them clearly. But it does not perform the evaluative act that jurists treat as constitutive of *'illa* identification. "Evaluative" here covers two different things. In one sense, it means weighing competing candidate causes against each other, a comparison among options the tradition has already articulated. But the deeper evaluative act is not a comparison at all. It is understanding what a ruling is *for*, its underlying wisdom (*ḥikma*), and reasoning from that understanding to a case the text does not explicitly settle. This distinction matters technically as well as philosophically. A model can perform well on cases that resemble configurations already present in its training data, since this is a matter of pattern-matching to precedent it has effectively seen before. The harder test is performance on a genuinely novel case, one whose relevant features do not resemble anything the model was trained on.⁵⁶ This is the test that actually matters for *ijtihād*. AI can be built to approximate the first kind of evaluation, weighing stated options against stated criteria. It has no way to perform the second, because there is no fixed list to search, only an act of understanding that has to be exercised fresh on unprecedented ground.

The prohibition of *ribā* in gold and silver illustrates how *'illa* identification unfolds across centuries. It illustrates, in particular, the second, deeper kind of evaluation just described. As I have examined elsewhere, Shāfi'ī jurists identified competing operative causes for the prohibition: gold and silver's function as monetary standards (*jawhariyyat al-thaman*), the value of things in general (*qiyam al-ashyā'*), or the value of consumables (*qiyam al-mutlafāt*).⁵⁷ Each formulation generates a different scope of analogy. Jurists initially described the *'illa* as inextensible (*qāṣira*) to gold and silver, but qualified this with "for the most part" (*ghāliban*), preserving the possibility of extension under changing economic conditions. That qualification only makes sense as an act of *ḥikma*-reasoning. The jurists were not simply choosing among a fixed set of candidate causes once and for all. They were preserving room to ask, as circumstances changed, what the prohibition was actually for, and whether that purpose still applied to a changed situation. The transition from metallic currencies to promissory notes and then to fiat money forced repeated

⁵⁶The study introduced in §1.2 above is worth returning to here. Salinas et al. found that law professors preferred LLM-generated short answers to contracts-law questions over their own colleagues' answers in a substantial majority of trials, with the strongest models outperforming all but one instructor. The result is compatible with the argument made here, not a counterexample to it, though not for the reason a first pass might suggest. The authors' own calibration analysis found that observed LLM win rates exceeded what stylistic features such as length and clarity predicted, so the preference is not simply an artifact of presentation. What the study measures, however, is a one-shot, forced-choice preference between two prepared answers to a settled doctrinal question. It does not measure whether an answer survives live, adversarial testing: al-Qarāfi's investigative questioning of a petitioner, the verification demanded by *taḥqīq al-manāṭ*, or cross-examination in a *qaḍā'* proceeding. Producing the better answer to a question already fixed in place is a different task from defending case-specific evidence against a real interlocutor's challenge, which is the competence this paper is concerned with.

⁵⁷Yousef Aly Wahb, "From Gold to Paper: The Applicability of Ribā to Modern Currencies in Shāfi'ī Jurisprudence," *UCLA Journal of Islamic and Near Eastern Law* 21, no. 1 (2024): 65-101.

reassessment across centuries. Fiat currency does not merely resemble a configuration already present in the classical corpus; it is a new economic reality the original jurists could not have anticipated. An AI system can reproduce the dominant view fluently, but that view is exactly what later jurists questioned. Determining whether modern currencies fall within the *'illa* requires the kind of contextual and evaluative judgment that classical jurists understood as *ijtihād*.

4.2 The Epistemic Problem: AI Cannot Verify Its Own Reasoning

In Islamic legal methodology, a valid ruling must be traceable to its justificatory basis, text, *'illa*, or recognized indicator. A ruling without traceable reasoning is merely assertion. By traceable reasoning, I mean the ability to identify the specific textual evidence, the operative cause, and the method of derivation, and to defend each under scholarly scrutiny. This is not merely the computational transparency of how an algorithm weighted its inputs. An LLM can produce plausible legal justifications, but the reasons it states are not the motivations for its output. It generates text through probabilistic continuation, not through inferential derivation from premises.

The problem of hallucination makes this limitation visible. When a model attributes a hadith incorrectly or cites a position not found in the tradition, the error is not accidental. It reflects the absence, in current generative language models, of an internal mechanism for verifying correspondence between output and source reality. Retrieval-augmented generation and citation-grounding pipelines have substantially reduced fabricated attributions, but these are external scaffolds bolted onto the system, not an internal *'adāla* and *dabt* of the transmitter. The Islamic tradition anchored reliability in trustworthy persons, documentation on its own was never sufficient. Even with these mitigations, systems trained exclusively on verified corpora can generate false statements, because the generation process is grounded in statistical likelihood. In a tradition that developed the *isnād* specifically to prevent unverified claims from entering the scholarly record, a system that generates unverifiable claims as a structural consequence of how it produces text is fundamentally at odds with Islamic epistemology.

Conclusion

This paper has argued that generative AI cannot independently perform the functions of *iftā'* or *qaḍā'*. The argument rests on three structural grounds: that the formalism/realism parallel reveals AI to be incapable of the contextual judgment both legal traditions treat as essential; that the classical conditions of the mufti and the judge are conditions of personhood that no algorithm can satisfy; and that even the technical operations of legal reasoning, from *'illa* identification to the verification of an AI system's own reasoning, require evaluative judgment at every step.

None of this denies that AI is a powerful tool. It can retrieve and organize existing rulings, analyze documents, structure case narratives, identify patterns across large datasets, and present competing scholarly positions with fluency. The American Arbitration Association's AI Arbitrator, launched in 2025 for documents-only construction disputes, illustrates the model: AI analyzes submissions and generates a structured case narrative, but a human arbitrator reviews the analysis, evaluates the outcome, and issues the binding award.⁵⁸ In every successful deployment of AI in legal contexts, the human retains the judgment function. AI assists in the administration of law but does not exercise the authority to interpret and apply it.

The foregoing analysis makes it possible to map, step by step, which operations in the *iftā'* and *qaḍā'* process AI can assist with and which it cannot.

The first step in *iftā'* involves parsing the questioner's words and assessing their condition, circumstances, and intent. While AI can receive and parse text, it cannot perform the assessment that gives the question its legal significance. A human mufti is required to read the question through the lens of *fiqh al-naḥs*.

The second step is gathering the relevant primary texts and evidences, where AI can be most helpful. It can search across digitized collections, identify relevant passages, and cross-reference sources with speed and comprehensiveness that exceeds what any individual scholar can achieve manually.

The third step is gathering precedent: identifying how earlier scholars ruled on the same or similar questions, mapping the positions of different schools, and presenting the range of existing opinions. As a retrieval task, AI can do this effectively. But these tasks are not merely preparatory labor on the way to a ruling. This is also how *malaka* is built in the first place. As al-Subkī's definition of the *mujtahid*, quoted above, already makes clear, that capacity is built only through repeated practice, and Bourdieu's habitus is formed the same way, through repeated immersion in exactly this kind of work. Handing this step to AI does not just let a scholar's faculty go soft. Done early enough, it can keep the faculty from forming in a novice at all, the way outsourcing every rep to a machine would.

The fourth step is identifying the operative cause and verifying its application to the case at hand. AI can present candidate *'ilhas* from the scholarly literature. The mufti evaluates which candidate applies to the particular case before him, as that evaluation depends on contextual, purposive, and often competing considerations that the formal rules do not determine. As the *ribā* case study demonstrated, this step has occupied jurists for centuries precisely because it resists mechanical resolution.

The fifth step is weighing competing opinions and selecting the strongest position. Where multiple valid opinions exist, the jurist must assess which is most appropriate given

⁵⁸ American Arbitration Association, "AI Arbitrator: Fast and Fair Dispute Resolution," accessed June 8, 2026, <https://www.adr.org/ai-arbitrator/>.

the questioner's context, and which is most consistent with the operative methodology of the school. *Tarjīh* is an evaluative act that depends on the jurist's understanding of the law's purposes and the case's particulars.

The sixth step is issuing the ruling. In *iftā'*, this is an act of *ikhbār* (reporting God's law for this case) that carries spiritual weight and scholarly accountability. In *qaḍā'*, it is an act of *inshā'* (origination) that creates a binding legal obligation. Both require a *mukallaf* who bears the consequences of the ruling before God and the community. This step is structurally impossible for AI, because the conditions under which a ruling is validly issued are conditions of the person issuing it, not conditions of the text produced.

AI, in short, can retrieve for steps two and three and surface candidates for steps four and five, but the evaluative judgment those last two steps require is not something it can supply. Steps one and six, which depend on *fiqh al-naḥs* and moral accountability, are closed entirely.

This mapping is a framework for responsible deployment. Muslims are already using AI tools for fiqh questions, and the real question is how AI will be used. The framework proposed here gives scholars and engineers a shared vocabulary for identifying where AI adds value, where it requires human oversight, and where it must defer entirely. It informs future collaboration on the design of tools that support the jurist's work without displacing it.

Two concerns remain even within the augmented model. The first is cognitive, and follows directly from the point made about step three above: reliance on AI tools, even as assistants, may erode the faculties this paper argues are essential. That concern is compounded by the observation that institutional *iftā'* in some contexts has already shifted toward retrieval-based operations, making the transition to AI-assisted processes appear seamless while further distancing the practice from its jurisprudential foundations.⁵⁹ The second is forward-looking. This paper's argument is grounded in the current state of AI, but it does not depend on current limitations alone. Even if artificial general intelligence were achieved, a system capable of human-level contextual reasoning, analogical thinking, and situated discernment, the question of *taklīf* would remain. Whether such a system could satisfy the conditions of moral agency, accountability before God, and the spiritual dimensions of the juristic faculty is a question that would require further theological and jurisprudential deliberation. But the argument made here bears on that question directly: capability is not the operative variable. Moral agency and accountability before God are conditions of personhood, not thresholds a system crosses by reasoning better, so even a hypothetically capable AGI would still need to be a moral subject, addressed by revelation and answerable for its rulings, before *taklīf* could apply to it. Whether any non-human system could ever be such a subject is the question that would require further theological and jurisprudential deliberation. This paper offers a framework for that conversation, not its conclusion.

⁵⁹ Tsourlaki, "Artificial Intelligence and Fatwā Issuance."