



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

Artificial Intelligence in Islamic Education and the Masjid

Curriculum, Language, Da‘wah, and Translation for Worship
and Education

Dr. Mohamed AbuTaleb

Dr. Mohamed AbuTaleb serves as VP of Research and Thought Leadership at Yaqeen Institute for Islamic Research, faculty and former Dean at several graduate seminaries, and Resident Scholar in the Research Triangle Region of North Carolina, USA. His research interests include Islamic institutional excellence, ethics of artificial intelligence (AI) & technology, and community leadership.

"الأراء في هذا البحث تعبر عن رأي الباحث وليس بالضرورة عن رأي أئما"

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




TABLE OF CONTENTS

1. Introduction.....	4
1.1 Current State of Discourse.....	5
2. The Framework in Brief.....	6
2.1 Underlying Technologies: Modes of Failure & Verification	6
2.2 Teachers and Transmission.....	7
3. Curriculum: Lesson Plans, Outlines, and Assessment	9
3.1 What the Instrument Does	9
3.2 Structure & Scheduling	9
3.3 Assessment	10
3.4 Advisable and Inadvisable Uses	11
4. Teaching Arabic & Correcting <i>Tajwīd</i>	13
4.1 Arabic Instruction	13
4.2 <i>Tajwīd</i> Evaluation: Current State Capabilities & Future Outlooks	13
4.3 Accessibility & Enablement.....	14
5. Da’wah Content and the Preparation of the Khuṭbah	16
5.1 <i>Amānat al-Minbar</i> and Trust of the Audience.....	16
5.2 Understanding Precedes Conveyance	16
5.3 The Risk of Eroding <i>Baraka</i> , Guidance, and Skill.....	17
5.4 Advisable and Inadvisable Uses	17
<i>Advisable Uses that Assist the Khaṭīb</i>	18
<i>Inadvisable Uses</i>	18
6. Real-Time Translation in the Khutba and Salāh.....	19
6.1 Scope and Fiqhi Context	19
6.2 Translating the Khutba	19
6.3 Translating Recitation of the Qur’ān.....	20
6.4 Translating Du‘ā’ and Qunūt	21
6.5 Advisable and Inadvisable Uses	22
7. Cross-Cutting Pitfalls	23
7.1 Failures in Disclosure and Documentation	23
7.2 Data Security & Privacy.....	23
7.3 Imitation of Scholarly Voice	23
7.4 Infrastructural Dependence & Digital Sovereignty	23
8. Conclusion	24

1. Introduction

Islamic ethics seek to proactively judge and shape the form, substance, and purpose of technology from its inception, expanding the purview of some other ethical frameworks that primarily mitigate harms after the fact. This is precisely the urgency of the questions examined here: in each of the four domains explored here in the realms of Islamic education and worship, the technology has already arrived and adoption is already underway. Imams draft *khuṭab* with these tools under the pressure of a Thursday evening; phone applications now correct makhārij for Muslims in their recitation of Quran; seminary teachers generate quizzes and lesson plans for their students; and congregants bring to technological systems the questions they once brought to their imam, including questions about grief, doubt, and the pivotal moments of their own lives¹. The fiqh and Islamic scholarship have not been silent to date and training has begun for community leaders and practitioners², though fertile ground remains for inquiry and research. The question is therefore not whether to permit practices that have yet to begin: it is how to discipline practices already underway.

In Dhū al-Qa'da 1446 / May 2025 the Council of the International Islamic Fiqh Academy, convened in Doha, issued Resolution 258 (3/26) on artificial intelligence. The Council held that the default ruling on its development and use to be permissibility, *ibāḥa*, conditioned upon six ethical and legal requirements.³ A number of research papers were submitted and examined in that session, and its deliberations advanced the position that there is necessary engagement and inquiry regarding this technology that falls upon the ummah. The broad responsibilities outlined include learning it, teaching it, and deriving benefit from it, in contrast to the risk of lagging behind and becoming dependent upon others. What the Resolution does not do, as befits a foundational ruling, is to specify based on domain. That is the principal objective of this work: to examine the applications such as the seminary classroom, the transmission of recitation, or the minbar.

That specification is the focus of this paper, which examines four domain applications: (1) curriculum, (2) the teaching of Classical Arabic and the correction of tajwīd, (3) da'wah content and the preparation of the khuṭbah, and (4) real-time translation in worship. A companion paper treats four domains of Islamic research — scholarly productivity, manuscript analysis, semantic search and takhrīj, and the limits of machine

¹ "AI Tools Are Transforming Muslim Worship. Religious Scholars Are Conflicted," Time, May 26, 2026, <https://time.com/article/2026/05/26/ai-muslim-worship/>.

² Waleed Kadous and Hidayath Ansari, "AI for Imams: For Imams and Those Active in Muslim Community Work" (workshop, hosted by AMJA Online), <https://cluesmith.com/ai-for-imams/>. Cited throughout as the most substantial practitioner-facing training resource presently available to imams in North America.

³ International Islamic Fiqh Academy (Organisation of Islamic Cooperation), Resolution No. 258 (3/26), "Artificial Intelligence: Its Rulings, Guidelines, and Ethics," twenty-sixth session, Doha, Dhū al-Qa'da 1446 / 4–8 May 2025, <https://iifa-aifi.org/en/56035.html>. The six conditions concern lawfulness of purpose and funding; realization of benefit and prevention of harm; non-disparagement of religion and religious symbols; protection of information and of public and private rights and freedoms; avoidance of threats to individual, societal, and national security; and adherence to honesty, proper documentation, and transparency in use.

translation — and shares the framework summarized in §2 below⁴. Both papers advance framework guidelines (*dawābīṭ*) rather than definitive legal rulings (*fatāwā*).

1.1 Current State of Discourse

The literature on Islamic AI ethics has moved well beyond first reactions, and the refusal of the optimism-pessimism binary is now the settled starting position of the field rather than a novel concept within it.⁵ One contribution stands closest to the argument of §5 below. Mohammed Gamal Abdelnour holds that engagement with AI must move from the "mind" to the heart (*qalb*), grounding ritual leadership in intention (*niyyah*), God-consciousness (*taqwā*), and reverence (*khushūʿ*) as elements presupposing a conscious believing subject. Abdelnour invokes the notion of cognitive outsourcing to argue that offloading the sacred activity of reflection is potentially spiritually corrosive.⁶

The two papers differ in the question asked, and the difference is worth stating precisely because it defines what this paper adds. Abdelnour addresses whether a machine could occupy the office of the imam, and answers that the suggestion is a category error. The question pursued here is the more pragmatic and less examined one: what happens to the *khaṭīb*, the teacher, and the reciter who do hold the office and use these instruments to discharge its responsibilities?

Two literatures external to Islamic studies also bear on what follows: research on (1) cognitive offloading and deskilling, and (2) empirical evaluation of religious content generation.⁷ The first, drawn from cognitive science and the study of automation, asks what becomes of a human faculty when the task that it previously exercised is handed off to a tool: *offloading* names the immediate delegation of mental effort, and *deskilling* the slower erosion of competence that can follow. The second is a small but growing body of work that tests what these systems actually generate when asked religious questions, scoring outputs both for doctrinal accuracy and for whether the sources they cite exist at all.

⁴ Mohamed AbuTaleb, "Artificial Intelligence as an Instrument of Islamic Research: Inquiry, Manuscripts, Verification, and Translation in the Work of the 'Ulamā'" (forthcoming).

⁵ Yaqub Chaudhary, "Islam and Artificial Intelligence," in *The Cambridge Companion to Religion and Artificial Intelligence*, ed. Beth Singler and Fraser Watts (Cambridge: Cambridge University Press, 2024), 109–28; Shoaib Ahmed Malik, "Artificial Intelligence and Islamic Thought: Two Distinctive Challenges," *Journal of Islamic and Muslim Studies* 8 (2024): 108–15; Yusuf Çelik, "Answering Divine Love: Human Distinctiveness in the Light of Islam and Artificial Superintelligence," *Sophia* 62 (2023): 679–96; Ezieddin Elmahjub, "Artificial Intelligence (AI) in Islamic Ethics: Towards Pluralist Ethical Benchmarking for AI," *Philosophy & Technology* 36, no. 4 (2023): 73; Amana Raquib et al., "Islamic Virtue-Based Ethics for Artificial Intelligence," *Discover Artificial Intelligence* 2, no. 1 (2022): 11.

⁶ Mohammed Gamal Abdelnour, "Artificial Intelligence and the Islamic Theology of Technology: From 'Means' to 'Meanings' and from 'Minds' to 'Hearts'," *Religions* 16, no. 6 (2025): 796, <https://doi.org/10.3390/rel16060796>.

⁷ On cognitive offloading and deskilling, see Michael Gerlich, "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking," *Societies* 15, no. 1 (2025): 6, reporting cognitive offloading strongly correlated with AI tool use and inversely correlated with critical-thinking scores; and Hao-Ping Lee et al., "The Impact of Generative AI on Critical Thinking," in *Proceedings of CHI 2025* (New York: ACM, 2025).

2. The Framework in Brief

The framework underlying both papers is set out at length elsewhere.⁸ The elements that govern what follows and are stated here in summary.

2.1 Underlying Technologies: Modes of Failure & Verification

Artificial intelligence across the four domains before us does not denote a single underlying technology, and the division is not immediately apparent without awareness of the technical underpinnings. Curriculum generation and sermon preparation involve the use of *large language models* (LLMs), *tajwīd* correction engages *automatic speech recognition*, and real-time translation involves *machine translation*. Each of these is a different underlying technology, using separate neural networks and distinct training data, which also fail in different ways. A large language model fabricates fluently and confidently; a recognition model mishears systematically; and a translation model resolves ambiguity silently. Further complexity arises from the fact that modern systems often combine them into single multimodal models.

A large language model (LLM) does not consult a source when it answers; its core architecture produces the text its training renders probable.⁹ It is an uncritical follower (*muqallid*) of its corpus in the parlance of Islamic scholarship: reproducing what it has absorbed, weighted by frequency rather than by strength of evidence, with no epistemological evidence (*dalīl*) steering that output from an *usuli* perspective and no capacity to independently examine the underlying reasoning. In other words, human engineers impose artificial algorithmic steering to suppress harmful outputs or favor certain directions. In our analogy, this is akin to the *muqallid* being heavily moderated by a censor as opposed to becoming an independent thinker (*mujtahid*).

What follows is that fabrication is intrinsic rather than incidental. The errors it introduces are citations that look right. When applied to religious queries, this architectural flaw becomes immediately apparent. Evaluation of both general-purpose and Islamically grounded assistants documents exactly this: phrases attributed to verses that do not contain them, quotations not locatable in an author's editions, reports for which no source exists.¹⁰ Retrieval grounding improves traceability but does not discharge the duty to check accuracy.¹¹

⁸ Mohamed AbuTaleb, "Reading Artificial Intelligence through the Islamic Sciences of Transmission" (forthcoming), where the argument compressed in §2 below is developed at length.

⁹ To fix the limitations of purely statistical generation, developers now connect LLMs to external systems through approaches like retrieval-augmented generation (RAG). While this does cause the system to consult a source, the assertion on the mechanism of language generation remains accurate.

¹⁰ Abdullah Mushtaq et al., "Can LLMs Write Faithfully? An Agent-Based Evaluation of LLM-generated Islamic Content," presented at the fifth Muslims in Machine Learning workshop, NeurIPS 2025, arXiv:2510.24438.

¹¹ Assessment of leading retrieval-grounded legal research tools found error rates above seventeen percent for two major products and above thirty-four percent for a third: Varun

Drawing a comparison to Hadith sciences, this duty to verify is the AI technological analog to the chain of transmission (*isnad*) not guaranteeing accuracy of the underlying report. However, base LLMs *cannot* provide an *isnad* at all because their weights are an opaque mathematical black box. It is only when external tools (like RAG) are bolted onto the LLM that it gains an incomplete, artificial *isnad*. Hadith scholars had the ability not only to examine the conclusions of past scholars, but to scrutinize the underlying data they relied on for those arguments; this is not fundamentally possible in the same extent or manner with generative AI.

In the tradition's own science of transmission error, it was expected that a mistaken text carried the indicators of its own corruption. Examples of this could be a garbled word, a linguistically impossible construction, or a name that no *rijāl* work contains.¹² Such transmission errors became known by various names indicating the issue, such as misreading of dotted letter (*taṣhīf*), distortion (*tahrīf*), or fabrication (*kadhib*). Generative AI fabrications do not carry such marks, as they are generated by a mechanism optimized to produce text that *looks* right. Essentially, they are errors that bear confidence and tone of accuracy (*ṣiḥḥa*). The consequence is that verification here cannot be negative — noticing that something looks amiss — but must be positive: confirming that the cited work exists and says what it is said to say. This principle governs each *ḍābiṭ* below: verification becomes mandatory by default rather than triggered by apparent defect.

2.2 Teachers and Transmission

The tradition held that certain knowledge is transmitted only through reception from a living and authorised teacher (*talaqqī*). The famous scholarly proverb states “He whose teacher is his book, his mistakes are more than his correct points”, warning against the one who takes his learning from the page alone. The figure of the *ṣaḥāfī*, or *muṣḥafī*, recurs in the literature of scholarly etiquette as a cautionary warning: one whose errors outweigh his correctness because no teacher stood between him and his misreading. Ibn Jamā‘a treats the relationship of teacher and student as a sanctified, spiritual bond designed for the cultivation of character rather than only mechanical information transfer. In other words, the learning process relies on the transmission of conduct and etiquette as much as knowledge itself.¹³

The warning here is not against books, but the substitution of a text for a relationship. Therefore, it speaks directly to any instrument offering the appearance of instruction without the correction, the *tarbiya*, and the chain that instruction from a

Magesh et al., "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools," *Journal of Empirical Legal Studies* 22, no. 2 (2025): 216–42. In medicine, an automated audit found that between half and nine-tenths of responses were not fully supported by the sources they themselves cited: Kevin Wu et al., "An Automated Framework for Assessing How Well LLMs Cite Relevant Medical References," *Nature Communications* 16 (2025): 3615.

¹² These errors could be caused by mistakes beyond *tahrīf*

¹³ Badr al-Dīn Ibn Jamā‘a, *Tadhkirat al-sāmi‘ wa-l-mutakallim fī adab al-‘ālim wa-l-muta‘allim*, ed. Muḥammad b. Mahdī al-‘Ajmī (Beirut: Dār al-Bashā‘ir al-Islāmiyya, 1433/2012), 87.

qualified teacher carries.¹⁴ Alongside it stands the *ijāza*, which locates the authority to transmit in a person, conferred by a person. It is not a certificate of information that is possessed; rather, it is the attestation by the teacher to the competence of the student. The key question therefore is not what information the instrument is authorised to produce, but what the learner is ready and authorized to engage.

The four sections that follow adopt a common structure for each domain area: an account of what the AI as an instrument does; the uses that are advisable; the uses that are inadvisable, with the characteristic pitfalls of the domain; and a proposed *dābiṭ*.

¹⁴ For a concise treatment of the role of scholars for the seeker of knowledge, see Usaama al-Azami, "Why We Need Ulama: The Importance of Seeking Islamic Knowledge from Scholars," Yaqeen Institute for Islamic Research, May 4, 2023, <https://doi.org/10.65061/YBGM4891>.

3. Curriculum: Lesson Plans, Outlines, and Assessment

3.1 What the Instrument Does

A curriculum is a plan for learning: an intentional selection and organization of aims and objectives, content, learning and teaching experiences, and means of evaluating outcomes.¹⁵ A lesson plan is the plan for a single instructional session within it. On the spectrum between the two, syllabi, course outlines, and unit plans occupy intermediate levels of scope and specificity. The curriculum represents the intended journey of learning, and the lesson plan is where it becomes enacted.

In articulating educational plans, AI appears to be unusually strong. Generative models can produce full lesson plans, including objectives, sequence, activities, checks for understanding, and assessment. Prompted to produce fifty examination questions on a topic, it will do so; asked to differentiate a lesson for three ability levels, it does so.¹⁶ Part of the technology's strength here is that it has absorbed many instances of these formulations, gleaning patterns from the collective wisdom and approaches of a myriad of teachers. This strength is also the source of its main hazard: a curriculum is not in fact a pure pattern-based formulation of the transmission of information, even if it superficially looks so.

3.2 Structure & Scheduling

In the Islamic educational tradition, the organization of a course of study is itself a scholarly achievement. The sequencing of material, experiences, and assessments encodes a theory of formulation that has been refined across generations. Further, curricula represent a tailored application to particular context, whether in age, societal conditions, language, or prerequisite knowledge. Each decision carries scholarly discretion and weight: why *naḥw* frequently precedes *uṣūl*, why a text is frequently memorized before it is analyzed, why certain works are withheld until a student has the prerequisite tools to read them competently, why *adab* is taught alongside content rather than after it. Indeed, the approach of an established school is its concise argument about *tarbiya*.

An AI model sequencing information has no access to those underlying arguments. It produces an ordering that is plausible, with topics arranged from simple to complex in the manner of the general educational literature it absorbed. It lacks an assumed understanding of the target learning audience and the corresponding methodologies required. Its output will look entirely reasonable to anyone who does not already know why the traditional ordering is constructed in a particular manner. This is the particular challenge: the failure mode is invisible to precisely the person most likely to rely upon the instrument.

A second and subtler drift follows from the typical training data, and its tendency to be grounded in a secular worldview, the most prominent Internet languages, and foreign

¹⁵ Hilda Taba, *Curriculum Development: Theory and Practice* (New York: Harcourt, Brace and World, 1962), 10–11.

¹⁶ On lesson plan generation, interactive quizzes, critique-and-improvement loops, and gamified learning as demonstrated to community educators, see Kadous and Ansari, "AI for Imams," session eight.

epistemologies in contradiction with the Islamic worldview.¹⁷ A generative AI model asked to produce Islamic curriculum will import the pedagogical assumptions of the corpus it learned from, which is overwhelmingly secular and shaped by these particular educational philosophies by default. The resulting drift is subtle, producing something harder to catch than outright doctrinal issues like *kufr* or *shirk*. Rather, a lesson could open a settled matter of *‘aqīda* as a topic for open student debate; an activity might invite learners to construct their own interpretation of a verse independent of any methodological bounds; a learning objective might promote a particular position on a matter of *fiqh* without awareness that other legitimate scholarly views exist within the mainstream. Each of these appears appropriate and even sophisticated to the unaware reader, and quietly shifts the formative methodology of thinking.

3.3 Assessment

Assessment deserves separate mention, because it is where the instrument is at once quite useful and yet potentially quietly corrosive. A generative AI model readily generates assessment question types that exhibit surface patterns, such as information recall, definition questions, or identification of rulings. It is less adept at judging deep cognitive processes, identifying genuine conceptual transfer, or assessing problem-solving skills. LLMs predict word patterns and lack genuine comprehension: they can not understand if a student actually understands how evidence supports a ruling or conclusion.¹⁸ This assessment gap was known in educational theory prior to the advent of generative AI. In the literature, it is named *construct underrepresentation*: the condition in which an examination samples only a portion of what it purports to measure, so that its results drive a narrower inference than the teacher intends.¹⁹ Modern educational practice has taxonomies distinguishing recall from reasoning, and established procedures for auditing whether the aggregate of a test's items matches the cognitive demand of the curriculum it serves.²⁰ Further, substantial evidence shows that examination can end up driving the teaching rather than reflecting it. Once a measure becomes a target it stops working as a metric, and whatever a course never tests, students learn to treat as unimportant.²¹

Generative AI neither created these issues nor is necessary to explain them; rather, it dramatically changes the cost-benefit analysis. AI effectively reduces the price of producing precisely those items that underrepresent the construct to nearly nothing, and

¹⁷ For an exploration of the impact of colonization on Islamic education systems, see Faisal Malik, "All That We Lost: The Colonized Mind and the Decline of the Islamic Education System," Yaqeen Institute for Islamic Research, August 15, 2019, <https://yaqeeninstitute.org/read/paper/the-colonized-mind-and-the-decline-of-the-islamic-education-system>.

¹⁸ Lydia Elliott, "Why AI Is Changing What It Means to Be Intelligent," Faculty Focus, July 31, 2026, <https://www.facultyfocus.com/articles/educational-assessment/why-ai-is-changing-what-it-means-to-be-intelligent/>.

¹⁹ Samuel Messick, "Validity," in *Educational Measurement*, 3rd ed., ed. Robert L. Linn (New York: American Council on Education and Macmillan, 1989), 13–103.

²⁰ Benjamin S. Bloom et al., *Taxonomy of Educational Objectives: The Classification of Educational Goals, Handbook I: Cognitive Domain* (New York: David McKay, 1956); Lorin W. Anderson and David R. Krathwohl, eds., *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives* (New York: Longman, 2001), among others.

²¹ George F. Madaus, "The Influence of Testing on the Curriculum," in *Critical Issues in Curriculum*, ed. Laurel N. Tanner (Chicago: University of Chicago Press, 1988).

raises correspondingly the educator's discipline required to refuse them. A teacher who accepts a generated item bank has made no error that any individual question would reveal. He has allowed his examination to migrate toward the testable and away from the core objectives of the discipline, so that it now measures only a subset of what he intended to teach.

This difficulty is even more pronounced still for objectives of *ta'lim* that do not submit to direct examination at all, such as the sincerity (*ikhlaṣ*), character (*adab*), and striving (*mujāhada*) of the student. Here the problem long precedes AI as well, and the literature that has treated it most seriously is not that of testing but of professional formation: the study of how clergy and physicians are certified in dispositions and judgment that no item can probe.²² The conclusion is that no single instrument can validly assess a complex competence, and that what is required instead is a program of many low-stakes observations aggregated by the judgment of a qualified assessor, terminating not in a single score but an act of entrustment.²³

The description here is quite familiar to Muslim scholars. This is essentially the *ijāza* system rediscovered in the context of modern professions: an attestation issued on the basis of extended companionship (*ṣuḥba*), encompassing both fidelity of transmission and strength of character. As a concrete example, in the sciences of *ḥadīth*, a transmitter's *'adāla* is established through corroborated testimony from those who knew him, precisely because integrity does not sit for an examination. Essentially, the Islamic tradition held a fitting instrument for these ends well before the modern discipline of assessment within educational theory succeeded in naming the problem. The risk in the AI age is not that these tools will fail to measure *adab*, but rather that the ease of generating what can be scored will make what cannot be scored appear optional. This directly motivates a discussion of advisable and inadvisable uses of the AI instrument in this domain.

3.4 Advisable and Inadvisable Uses

AI should elaborate on sequencing and structure that the teacher has set rather than set the sequence itself. A teacher who supplies his own interpretive approach and the text he is teaching, and asks for activities, worked examples, practice exercises, and intermediate checks for understanding retains the formative judgement where it belongs. Further, AI can be helpful in formulating differentiated materials for mixed-ability classes, which is laborious and repetitive work. The instrument can also draft question banks that the teacher vets, re-weights, and supplements with the reasoning items the model will not produce; and can absorb the substantial administrative burden of teaching, such as formulating rubrics, schedules, and progress summaries, which all consumes time that could be spent on instruction and do not carry formative weight for the student's learning.

²² Lee S. Shulman, "Signature Pedagogies in the Professions," *Daedalus* 134, no. 3 (2005): 52–59; Charles R. Foster et al., *Educating Clergy: Teaching Practices and Pastoral Imagination* (San Francisco: Jossey-Bass, 2006).

²³ Cees P. M. van der Vleuten and Lambert W. T. Schuwirth, "Assessing Professional Competence: From Methods to Programmes," *Medical Education* 39, no. 3 (2005): 309–17; Olle ten Cate, "Nuts and Bolts of Entrustable Professional Activities," *Journal of Graduate Medical Education* 5, no. 1 (2013): 157–58; Brian D. Hodges, "Assessment in the Post-Psychometric Era: Learning to Love the Subjective and Collective," *Medical Teacher* 35, no. 7 (2013): 564–68.

Inadvisable uses include the generation of a curriculum from scratch by anyone not already competent to evaluate its ordering; the adoption of generated assessment without deliberate attention to what is being measured; and the use of generated materials on matters of *'aqīda* or contested *fiqh* without review by someone qualified to catch the drift described above, which is the class of error least likely to be perceptible by a non-specialist in the field.

Proposed Guideline (dābiṭ): The instrument may elaborate upon a curriculum, but should not independently sequence one without specialized review. Curricular tartīb encodes tarbiya, and the formative aim and selection of assessment must remain with the teacher.

4. Teaching Arabic & Correcting *Tajwīd*

As two substantially different AI technologies underlie these applications, they should be treated apart.

4.1 Arabic Instruction

For language teaching the relevant instrument is the large language model (LLM) and it offers several compelling strengths in this domain. AI competently generates drills for declension (*i'rāb*) and morphology (*ṣarf*), teaches vocabulary in context, formulates example sentences, and provides adaptive, real-time practice. The AI instrument also offers inexhaustible, patient repetition without embarrassment, removing a barrier for students and self-conscious adult learners in particular.

An important caveat to note for any teacher who intends to rely on these tools: models can make Arabic errors. Training corpora are heavily weighted toward Modern Standard Arabic and contemporary usage; the classical register is comparatively thin, and grammatical analysis is generated with the same confident fluency as everything else. A model will supply a wrong *i'rāb* in a confident tone. Further, while Arabic is a phonetic language with predictable word pronunciation and morphological patterns, its vocabulary is enormous and patterns of eloquence inexhaustible; thus, a particular usage may fall outside the dominant norm in classical poetry or historical usage. If this is true in human usage, it is even more of a concern regarding fidelity for the incomparable expression of the Holy Qur'ān. For drilling forms the instrument is dependable; for adjudicating a grammatical analysis it is not universally so, and a student without a teacher has no way to know which he has received.

4.2 *Tajwīd* Evaluation: Current State Capabilities & Future Outlooks

For recitation correction, the underlying technology is automatic speech recognition. Evaluation of its accuracy requires nuance, because the published figures are both impressive and narrower than they appear. Accuracies in the high nineties are common, but they are typically obtained on classifiers trained to detect four or eight specific rules across small purpose-built datasets, not comprehensive *tajwīd* detection across the entire Qur'ān.²⁴ Perhaps more instructive is a study on Juz' 30 that reported roughly 91% accuracy while acknowledging that it remained unable to distinguish certain Arabic letters of similar pronunciation.²⁵ That admission goes to the heart of the matter, because the distinctions in question are those of *mutashābih al-makhārij* — those letters sharing similar points of vocal

²⁴ Reported figures include ninety-nine percent accuracy on four *tajwīd* rules across roughly six hundred and fifty records, and approximately ninety-six percent across eight rules on some three thousand single-rule audio files. See "Mispronunciation Detection of Basic Quranic Recitation Rules using Deep Learning," arXiv:2305.06429, and "Speech Recognition Models for Holy Quran Recitation Based on Modern Approaches and Tajweed Rules: A Comprehensive Overview," International Journal of Advanced Computer Science and Applications 14, no. 12 (2023). These are single-rule classifiers on small purpose-built corpora, not general assessments of recitation.

²⁵ "Speech Recognition Models for Holy Quran Recitation Based on Modern Approaches and Tajweed Rules: A Comprehensive Overview," International Journal of Advanced Computer Science and Applications 14, no. 12 (2023).

articulation — and the capacity to hear the difference between letters that resemble one another is a great deal of what *tajwīd* competence consists of. Mispronunciation of letters is considered a serious error that can affect a change in meaning entirely.

It would nonetheless be a mistake to build an argument on present inadequacy alone. The trajectory of improvement is steep: recent work has trained on more than 850 hours of expert recitation using a phonetic script encoding, incorporating not merely phonemes but the *ṣifāt* of articulation, reporting phoneme error rates well below 1%.²⁶ It may be that within a short period these systems will detect articulation errors as reliably as a competent human listener, and possibly even more so. The objection worth noting is therefore not that the instrument is inaccurate: it is that accuracy is not the whole of what a teacher of Qur'ān provides.

A teacher of recitation performs several separable functions. Critically, the teacher imparts *adab*, meanings, and gems from the Qur'ān, which we have noted in the previous section cannot be accomplished by the AI instrument. The teacher diagnoses and corrects errors, coaching the student through practice and providing feedback. Perhaps most importantly, he transmits the recitation as he received it, standing in a chain that reaches back through the Prophet Muhammad (saw) to Angel *Jibreel* to Allah (swt) Himself; and attesting that his student has received it from him after certifying the requisite fidelity. The matter of accuracy is regarding *ṣiḥḥat al-nuṭq* (proper pronunciation) and is increasingly within reach of a machine. The matters of cultivating character and of *ittiṣāl al-sanad* (connectedness of the transmission chain) are permanently out of reach of AI as an instrument. An *ijāza* in Qur'ān certifies reception from a named person who himself received from others: no model has received anything, and none holds an *ijāza* to confer.

It follows then that the machine's proper place is one the tradition already accommodated before its invention. Students have always practiced alone, drilled with peers, and corrected themselves between sessions before presenting to the shaykh. An application that corrects *makhārij* between lessons may ultimately become the technically superior form of that practice, and there is no reason for hesitation in incorporating it should it reach that fidelity. What it cannot do is stand at the end of the process and attest, nor shape the embodiment and understanding of the Quran. To be clear, the presentation to a qualified teacher who holds a chain is not a formality that a sufficiently accurate machine renders redundant; it is the very substance of the transmission.

4.3 Accessibility & Enablement

Our consideration is not restricted to capability or caution alone, but also recognizes benefits that are unlocked by this technology. For example, one recitation application now facilitates Qur'ān engagement at a staggering scale — on the order of millions of hours in a single Ramadan across more than a hundred and eighty countries.²⁷ For a Muslim in a town with no qualified teacher, for a convert with no community, for a woman who cannot readily access instruction, the relevant comparison is not between a machine and a shaykh. It is between a machine and nothing at all. The *muwāzana* in such a case is not close.

²⁶ Abdullah Abdelfattah, Mahmoud I. Khalil, and Hazem Abbas, "Automatic Pronunciation Error Detection and Correction of the Holy Quran's Learners Using Deep Learning," preprint, submitted August 27, 2025, arXiv:2509.00094, <https://doi.org/10.48550/arXiv.2509.00094>.

²⁷ See "AI Tools Are Transforming Muslim Worship," Time, May 26, 2026.

Proposed Guideline (qābiṭ): Given sufficient accuracy benchmarks, the AI instrument may diagnose the technical recitation; it cannot independently teach adab, embodiment, or meaning. It best serves as a subordinate, supplementary aid (mudarris mu'īn) between sessions and as a means of access where no teacher is available, but no machine assessment substitutes for presentation to one who holds a chain, and no technology application confers an ijāza.

5. Da'wah Content and the Preparation of the Khuṭbah

5.1 *Amānat al-Minbar* and Trust of the Audience

Each *khaṭīb* who ascends the *minbar*, convention speaker who assumes the lectern, activist who addresses an assembled crowd, or speaker who stands behind the camera lens feeding an online dawah video takes on a sacred trust from their audience. With AI entering the process, Resolution 258 supplies ethical grounds for consideration with the requirement of honesty and transparency, which we move to apply specifically to this domain. When a *khaṭīb*, speaker, or preacher addresses an audience, those before him implicitly understand him to be conveying what he has personally understood and represents as true. That presumption is not defeated by his having referred to books, consulted scholars, or used peers as a sounding board in preparation — no one supposes a *khaṭīb*'s material to be entirely original to him, and no one expects him to enumerate his sources as part of the sermon. However, this implicit understanding is defeated should the speaker deliver words that have not passed through his own comprehension. The dividing line therefore does not fall at the origin of the material, but at whether the material was understood and owned before it was spoken.

In this manner, the role of generative AI maps naturally to this pre-existing framework. A *khaṭīb* who uses AI to locate evidences, to structure an argument, or to draft a passage he then studies, revises, and comes to hold has done nothing requiring explicit disclosure, because the presumption of the *minbar* remains satisfied: the words are his, and he takes ownership and responsibility of the ideas conveyed. A *khaṭīb* who reads out text he encountered for the first time an hour earlier and has not thought through has broken the trust of his congregation, whether or not a machine was involved.

The tradition supplies a framing necessary to differentiate advisable from inadvisable uses of AI, which we describe in the following section.

5.2 Understanding Precedes Conveyance

The Qur'ān describes the responsibility of religious scholarship in a verse in Surat at-Tawba, offering several actions in a particular sequence. A party goes forth from among the believers "that they may acquire understanding in the religion, and that they may warn their own people when they return to them, so that they may beware."²⁸ Ibn Ashur points out that *li-yatafaqqahu* implies a rigorous, strenuous effort to acquire deep understanding.²⁹ as-Sa'dī emphasizes that the goal of seeking Islamic knowledge is societal benefit: the scholar must integrate back into his community and actively teach, counsel, and warn them.³⁰ The order here is instructive and important: going forth and acquiring understanding precede addressing the people.

²⁸ Qur'ān 9:122.

²⁹ Muḥammad al-Ṭāhir ibn 'Āshūr, *Tafsīr al-Taḥrīr wa-l-tanwīr* (Tunis: al-Dār al-Tūnisiyya li-l-Nashr, 1984), 11:63, ad Q 9:122.

³⁰ 'Abd al-Raḥmān ibn Nāṣir al-Sa'dī, *Taysīr al-Karīm al-Raḥmān fī tafsīr kalām al-Mannān*, ed. 'Abd al-Raḥmān ibn Mu'allā al-Luwayḥīq (Beirut: Mu'assasat al-Risāla, 1420/2000), 357, ad Q 9:122.

This yields the distinction the present question requires. The means that serve khutba preparation are many, and the AI instrument offers new efficiencies and possibilities in the means of preparation. Locating a ḥadīth, verifying a citation, arranging material in a defensible order, translating a passage, anticipating the objection a listener might raise, producing a follow-up study sheet for the congregation — these activities all predate AI. The core objective and precursor to delivery is the *tafaqquh* itself, and *tafaqquh* cannot be outsourced or delegated: it is a state acquired in a person.

The failure case can now be described precisely, and without any claims regarding sincerity or inspection of the *khaṭīb's* heart. He who ascends the minbar with text he has not understood stands in the position of one who returns from a journey he did not take. He has retained the *indhār* and skipped the *tafaqquh*, and the verse makes the former the successor of the latter.

The verse also helps us understand another key factor in its description of the addressees as “*their own people*.” One cannot warn a people without knowing who they are: their context, what has happened among them, what they are avoiding and what they are ready for. A model can generate counsel, but it lacks particular context to tailor the message to a particular people with Prophetic wisdom. The model does not stand in the relationship to them invoked by the verse.

5.3 The Risk of Eroding *Baraka*, Guidance, and Skill

There is a further consideration which bears upon the *khaṭīb* rather than the congregation. The Qur’ān states, “*As for those who struggle in Our cause, We will surely guide them along Our Way. And Allah is certainly with the good-doers.*”³¹ al-Baghawi cited a famous principle circulated among earlier scholars, “whomever acts upon what he knows, Allah will grant him the knowledge of what he does not know.” Further, Fakhr al-Din ar-Razi asserted that striving precedes guidance. Among what is gleaned from the verse is that the paths are opened to the one who exerted himself in seeking them, and that pursuit should be solely and sincerely for the sake of Allah.

In the case of deficiency in this dimension, the congregation may still receive a beneficial sermon; it is in fact the *khaṭīb* who is the one at loss. He risks producing an output without receiving the guidance promised to one who strove for it. A person who is no longer being guided through his own preparation has, over months and years, progressively less to give. Even if the sermons do not perceptibly deteriorate in form, he himself has begun to derive less transformative benefit than was once afforded by his striving. What one does not receive, one cannot transmit.

5.4 Advisable and Inadvisable Uses

The preceding discussion yields a workable test, and it is applied below to several present uses. A use is sound where it assists the going forth or serves the return; it is unsound where it performs the *tafaqquh* in the *khaṭīb's* place or inappropriately stands in relation to the audience that should be occupied by the *khaṭīb*.

³¹ Qur’ān 29:69.

Advisable Uses that Assist the *Khaṭīb*

Information retrieval and verification: locating an evidence, confirming a citation against its source, checking a *ḥadīth* authenticity; caution is warranted because hallucinated material now imitates accuracy, so verification must be proactive, consistent, and against the source itself rather than judged by the output looking sound.

Anticipating objections: asking for the strongest objection a listener might raise, the weakest link in his reasoning, or the ways a claim could be misheard.

Structural work: reordering, tightening, cutting to the time available.

Language: translating a passage he then checks, or assisting a *khaṭīb* who must preach in a language not his own.

Supplementary materials: study sheets, follow-up notes, an outline for a subsequent circle, material for those who were absent.

Rehearsal: explaining his argument and seeing whether it comes back to him intelligibly, using AI as a diagnostic.

Inadvisable Uses

Khuṭbah generation without review and independent thought. This is the main use case discussed above, and its risks have been outlined.

Asking what to speak about: this is problematic in absence of supplying additional context and exercising independent critical thinking. The choice of subject is the judgement the *khaṭīb* must exercise, and the AI instrument can assist at most.

Borrowed voice: producing a *khuṭbah* in the manner of a known scholar other than oneself, which takes away authenticity and resonance.

Preaching an unverified ruling: carrying to the minbar a position taken from AI and not confirmed with a qualified source.

Proposed Guideline (ḍābiṭ): *The AI instrument may serve the preparation of the khuṭbah, but the striving and understanding required for delivery must not be delegated. Nothing should reach the minbar that the khaṭīb has not himself understood, and the tailored, contextualized naṣīḥa that he owes to this particular audience.*

6. Real-Time Translation in the Khutba and Salāh

6.1 Scope and Fiqhi Context

This matter at hand is in fact comprised of three distinct questions, distinguished by what is being translated and in what setting: the instructional content of the *khutbah*; the recitation of the Qur'ān, whether within the sermon itself or in *ṣalāh*; and *du'ā'* and *qunūt*. Further, we note that fiqh councils have addressed matters of technologically mediated worship prior to the advent of generative AI. A noted example is specifically cited in the Council of the International Islamic Fiqh Academy's Resolution 258 on AI: its own ruling on *ṣalāh* performed behind a transmission by telephone, radio, or television.

Contemporary *fatāwa* have stated that worshippers who lack proficiency in Arabic may follow the imam's recitation during Taraweeh and other supererogatory prayers by reading a translation of the Qur'an's meaning, whether from a printed text or an electronic device.³² They have also stated there is no objection to using screens, monitors, or other devices to provide real-time translation of the *khutbah* for those who do not understand Arabic; or that earbuds may be used to hear a translation of the verses recited in Taraweeh or of the Friday *khutbah* itself.³³

6.2 Translating the Khutba

In the case of the explanatory and admonitory content of the sermon, the case for live translation is strong. The *khutbah* has an instructional function, and a congregation that cannot understand is rewarded for the performance of worship and gains the *barakah* of the gathering, but does not comprehend direct admonition or instruction. In a community of mixed language capabilities, typically including converts, students, non-Arabic speakers, and multiple countries of origin, the alternative to imperfect comprehension is often no comprehension at all.

The underlying standard driving the *fatāwa* is need (*ḥāja*), and that does not depend on the claim that translation is perfect or even excellent. It is sufficient to show that the need is satisfied by a means that delivers the benefit sought — here, access to the substance of an address otherwise opaque to the listener. Measured against that standard, the present state of machine translation is adequate: recent studies confirm that the broad themes of an address survive machine translation even where finer terminology does not.³⁴ These shortcomings appear in areas like precision in religious terminology, the speaker's

³² Assembly of Muslim Jurists of America, "Recommendations of the Conference on Contemporary Issues Related to New Muslims," under "Issues Related to Acts of Worship," AMJA Online, accessed August 3, 2026, <https://www.amjaonline.org/declaration-articles/recommendations-of-the-conference-on-contemporary-issues-related-to-new-muslims>.

³³ Assembly of Muslim Jurists of America, Recommendations of the 18th Annual Imams' Conference: Contemporary Issues Related to Technology (AMJA, 2022), <https://www.amjaonline.org/wp-content/uploads/2023/03/Recommendations-of-the-18th-Annual-Imams%E2%80%99-Conf-Contemporary-Issues-Related-to-Technology-Formatted.pdf>.

³⁴ Amal Abdelsattar Metwally, Wagdi Rashad Ali Bin-Hady, and Eisa Asiri, "Assessing AI-Powered Translation Quality: Insights from the Translation of the Farewell Sermon," *International Journal of Language Education* 8, no. 4 (2024), <https://doi.org/10.26858/ijole.v8i4.70034>.

stance and authority, and cohesion across the discourse as a whole. They are further aggravated by the latency constraints that live rendering imposes, since simultaneous systems operate with less context than asynchronous ones by design.

The practical implication of these technical shortcomings is therefore a question of deployment rather than of permissibility. Human mediated translation is also imperfect and often expensive, and a congregant who would otherwise understand nothing is better served by an imperfect rendering true to broad themes than by none at all. The same technical limitations, however, drive against treating machine output as a substitute for a prepared translation where one is available, and for the distinct treatment of Qur'anic recitation as advocated below. Further, it is important to incorporate human verification wherever translated content will be relied upon as instruction or reference rather than merely followed in the moment. A printed translation can be reviewed before publication, and a live translation should not be automatically posted to the Internet without review. Finally, disclosure is important: the congregation should know that what they are reading is machine-produced, and that AI can make mistakes. Verses of the Qur'ān recited within the sermon fall under the treatment of the following section rather than this one.

6.3 Translating Recitation of the Qur'ān

The Qur'an is a definitive and fixed text: its wording is fully known, transmitted by mass narration, and identical in every copy. This changes the technical problem entirely. A device helping a worshipper follow the imam does not need to translate anything at all; it needs only to recognize which verse is being recited and then display a translation that already exists.

That distinction is significant because the two approaches fail in different ways. A system that generates a translation as it listens can produce output that is fluent, confident, and wrong. The stakes are higher in this application than the other two. The worshipper it serves is, by definition, someone who cannot check it — that is why he is reading a translation in the first place. A system that identifies the verse and retrieves a published translation fails differently. Its worst error is displaying the wrong verse, and that error is visible: the text will not fit the flow of the recitation, and the worshipper has a better chance of noticing that something has slipped. One approach can put invented meaning in front of him without warning, whereas the other can only put a correct meaning in the wrong place.

Retrieval also brings accountability that generation cannot. A published translation has named translators, an editorial process, and an institution that stands behind it. A masjid can review it once and know exactly what its congregation is reading every night. Generated text cannot be reviewed in advance, because it does not exist until the moment it appears on the screen. Retrieval carries a further practical benefit: the same verse is rendered the same way each time, matching the printed copies congregants may already own and reinforcing rather than competing with what they are learning.

There is a deeper reason for the same conclusion. Scholars have long held that what is rendered into another language is not the Qur'an itself but a translation of its meanings — an interpretive act closer to *tafsīr* than to transcription. Whatever latitude may exist for machine translation elsewhere, interpretation of revelation is not a task to delegate to an unsupervised system. The distinction is the familiar one between transmitting what qualified scholars have already established (*naql*) and deriving meaning independently (*istinbāt*). A retrieval system transmits, whereas a generative system derives. Only the first is appropriate in this application.

A technical consideration supports the same conclusion. Quranic recitation is not ordinary speech: *tajwīd* lengthens vowels, shapes pauses, and carries unique intonation and nasalization, all of which degrade speech recognition systems trained on conversational audio.³⁵ Recognition of recitation is also a more constrained problem than open-domain transcription, since the set of admissible outputs is fixed and known in advance.³⁶

A masjid adopting this technology for Qur'anic recitation should require that the translation be drawn from an approved, human-reviewed translation, located by matching the recitation against the Qur'anic corpus. It should not deploy systems that generate Qur'anic translation in real-time.

6.4 Translating Du‘ā’ and Qunūt

Supplication is different from Quranic recitation in that it is not drawn from a fixed text. The imam may use transmitted narrations, his own words, or a mixture of the two. For the portions he composes himself there is no corpus to match against, and thus generative translation is required. In this setting, that limitation is tolerable, because what translation needs to accomplish is different. The congregant is not being taught a text he must learn precisely; he is being put in a position to know what is being asked of his Lord, so that he can raise his hands and say *āmīn* with presence of heart. What he needs is the substance of the request, such as a *du‘ā’* for forgiveness, relief for those suffering, protection, or rain. The congregant is not in need of the imam's exact phrasing in order to appreciate the meaning.

These requirements align closely with the technical capabilities of current systems, as discussed in the earlier section §6.2 on real-time translations of sermons. The specific gaps that are serious deficiencies in real-time translation of a Qur'anic verse or particular legal ruling are tolerable here, as what is produced by translation is what is needed. Further, the benefit here is important to recognize. A long supplication in a language one does not understand can invite the mind to drift, whereas understanding of the same supplication can invite humility, even tears. The purpose of the technology here is unlocking a higher level of participation for the worshipper to join in what is being asked of his Lord with comprehension.

Qunūt introduces one further consideration, which is that it takes place within the prayer rather than alongside it; the *fatāwā* cited above already address this, permitting a translation to be read from a printed text or an electronic device during *Tarāwīḥ* and other supererogatory prayers, within which scope *qunūt* in *witr* falls. It is also the most textually mixed of the three cases, since imams commonly join transmitted supplication, Qur'anic verses, and their own words in a single *qunūt*.

The safest, most accurate approach here is a hybridized one. Where the imam uses a well-known transmitted supplication, a system can recognize it and supply the established translation, reserving generated output for the imam's own words. The technical design

³⁵ Ali M. Alagrami and Maged M. Eljazzar, "SMARTAJWEED: Automatic Recognition of Arabic Quranic Recitation Rules," in *Computer Science & Information Technology*, ed. David C. Wyld et al. (CS & IT-CSCP, 2020), 145–52, <https://doi.org/10.5121/csit.2020.101812>.

³⁶ Wei Zhou et al., "RASR2: The RWTH ASR Toolkit for Generic Sequence-to-Sequence Speech Recognition," in *Interspeech 2023* (Dublin: ISCA, 2023), 4094–98, <https://doi.org/10.21437/Interspeech.2023-1062>.

principle here is a hybrid, retrieval-augmented approach: the system prefers retrieval wherever possible, and falls back on generation only where it is not.

6.5 Advisable and Inadvisable Uses

The three cases require tailored and differentiated approaches. The distinction between them is not the sanctity of the words alone but whether a reviewed translation of them already exists. Where the spoken words are recognizable from known narrations, recognition and retrieval replace translation, and the verification is complete before any worshipper reads a screen. Where speech is original to the speaker, no such translation can exist, and fidelity to substance is the most that can be asked. The principle across this section is accordingly to prefer retrieval wherever a reviewed translation can be matched to a known text, and to accept generation only where the source is not fixed and the aim is comprehension of general meaning.

Advisable uses therefore include live rendering of the explanatory and admonitory content of the *khuṭbah*; retrieval-based display of an approved translation for Qur'ān recited in *ṣalāh*, and equally for verses the *khaṭīb* recites within the sermon itself; recognition of well-known transmitted supplications with supply of their established translations; and generative rendering of the imam's own *du'ā'*, where the worshipper's purpose is to know what is being asked of his Lord so that he may assent to it with presence of heart. A hybrid system that draws on retrieval wherever a text can be matched and falls back on generation only where it cannot is the technically and juristically preferable design in each of these settings.

Inadvisable uses include the generation of Qur'anic translation in real time, which is the central restriction of this section; the treatment of live machine output as a substitute for a prepared translation where one is available, whether taught in *ḥalaqāt* or distributed in print; the republication of a live rendering, to the internet or otherwise, without human review, since material relied upon as instruction or reference stands on different footing than material followed in the moment; and the delivery of any machine rendering without informing the congregation of its nature.

*Proposed Guideline (qābiṭ): Real-time translation may serve the instructional content of the *khuṭbah* and the imam's own *du'ā'*, where the governing standard is need (*ḥāja*) and fidelity to broad meaning is sufficient to satisfy it. It should not generate the meaning of recited Qur'ān in the moment: translation of revelation is an interpretive act, and what is displayed must be transmitted (*naql*) from an approved, human-reviewed translation located by matching the recitation against the corpus. Transmitted supplications should likewise draw on established translations wherever they can be recognized. In every case the congregation should be informed that the rendering is machine-produced, and no live rendering should be republished or relied upon as instruction without human review.*

7. Cross-Cutting Pitfalls

Four pitfalls cut across all of the areas we have examined in Islamic education, as follows.

7.1 Failures in Disclosure and Documentation

Gaps in appropriate disclosure and sourcing are more subtle and difficult to detect than ever before, running counter to Islam's rich tradition of traceability and sourcing. Exacerbating the matter is an absence of agreed-upon standards in attribution and disclosure in the generative AI age. Examples include machine translation circulated without attribution; unreviewed generated teaching material presented as the institution's own; a sermon delivered as one's own understanding when it is not. Each is a breach of the sixth condition of Resolution 258 and easy to commit inadvertently. Standards are needed in order to expect adherence to common conventions.

7.2 Data Security & Privacy

As AI systems become more intertwined with education, research, and administrative processes, special attention is needed towards the matter of the sensitive data these systems ingest. In the course of processing AI requests, information submitted to a commercial system has been shared with a third party. People often trust their muftis, imams, and teachers with some of their most private and sensitive matters, such as marital difficulty, financial challenges, sin, or doubt. A questioner who approaches his imam does not expect that his circumstances will be typed into an external service, and standard commercial terms, particularly on free services, rarely afford the protection such material warrants. Resolution 258's fourth condition speaks directly to this. Institutional policies need to be developed and adopted to govern the acceptable use of these tools, and offer contextualized guidance on practical matters such as anonymization before submission, restriction of sensitive matters, and data retention policies, among others.

7.3 Imitation of Scholarly Voice

The AI instrument furnishes anyone with the means to fabricate a scholar's speech in convincing tone, proper grammar, and flowing rhetoric. In the present state, a convincing audio clip attributed to a known *'ālim*, expressing a ruling he never gave or a speech he never delivered, is now trivial to produce; sadly, the Internet is already riddled with such deepfakes. The office of the scholar is directly exposed to this risk. Communities need to establish how an attribution is authenticated and where a scholar's positions are authoritatively published.

7.4 Infrastructural Dependence & Digital Sovereignty

The systems described in this paper are, with a few notable exceptions, built and controlled elsewhere. These systems are trained on corpora the ummah did not curate, and governed by values it was not consulted upon. The argument that the Muslim community is repeating with artificial intelligence a familiar pattern of delayed adoption and consequent dependence converges with the collective obligation invoked in §1: to learn, to teach, and to build, rather than only to consume.

8. Conclusion

The four domains examined here share a commonality: in each, the AI instrument offers particular uses that are genuinely enabling, and others that are inadvisable or ask the tool to perform what is ill-suited to execute. None of these guidelines rests upon the present inadequacy of the technology, and this bears emphasis because arguments built on current limitations expire as the field advances. The trajectory of capability is steep in every domain surveyed, and to the best of our understanding, the guidelines proposed here would stand unchanged were the systems twice as accurate as they are today. What is reserved to the human is not reserved because a machine performs it poorly, but because *ittiṣāl al-sanad*, *tarbiya*, entrustment, and understanding are constituted by the presence of a person and are not the sort of thing an instrument can supply at any level of fidelity.

What the technology does change materially is the discipline demanded of the practitioner. Fabrication that bears the tone of *ṣiḥḥa* and carries none of the marks of *taṣḥīf* converts verification from a response to apparent defect into a standing obligation; assessment that is cheap to generate makes what cannot be scored appear optional; and a sermon that arrives complete in moments hides the loss of striving to which guidance was promised. In each case, the instrument does not create the failure so much as lower its price, and the burden falls upon the teacher, the reciter, and the *khaṭīb* to refuse what is now easy.

The conclusion is therefore not abstention. Particularly in contexts where qualified practitioners are few or inaccessible, the collective obligation identified in §1 — to learn, to teach, to derive benefit, and to build rather than merely to consume — is not discharged by caution. It is discharged by adoption that is disciplined, disclosed, and accountable to those it serves. The guidelines offered here are *qawābiṭ* and not *fatāwā*, advanced in the expectation that practice in these four domains may ultimately outpace them and necessitate their revision. Further, they serve as a starting point for institutions to apply to their own contexts in crafting policies to guide the execution of their missions in the generative AI age.