

The Trust We Cannot Delegate

Artificial Intelligence, Accountability, and the Human Mind

IDRIS NAWAZ

Futūḥ
Journal of Islamic Research by Miftaah Institute

Abstract

Artificial intelligence now writes, diagnoses, plans, and increasingly acts on our behalf, and with each advance the temptation grows to let it think for us. This article argues, from within the Islamic tradition, that such a surrender is not ours to make. It first describes in plain terms what contemporary AI can do, conceding its capability honestly. It then shows why, however capable it becomes, AI cannot bear *taklīf*, the moral and legal accountability that the Qurʾān places upon the human being alone. Because AI can only ever be an instrument (*sabab*), responsibility for whatever is done through it returns to us. That responsibility presupposes a living, exercised intellect (*ʿaql*). Drawing on recent research on cognitive offloading alongside al-Ghazālī's account of gratitude, the article argues that the habitual delegation of thought erodes the very faculty upon which our accountability rests. It closes with practical guidance for the individual and a set of goals for the Muslim community, concluding that AI may organize and inform human judgment, but must never replace it.

Keywords: taklīf, sabab, artificial intelligence, cognitive offloading, ʿaql

Introduction

We have built machines that can do almost anything we ask of them, except answer for what they do.

A new phrase has entered ordinary speech. When an email goes out with an error no one caught, when a report cites a source that does not exist, when an automated decision harms someone, the explanation increasingly offered is a simple one: the AI did it. The phrase sounds like a description, but it works as an excuse. It quietly moves responsibility from the person who used a tool to the tool itself, as though the system were a party that could be blamed, corrected, or made to answer. And the more capable these systems become, the more natural the excuse sounds.

The capability is real. Artificial intelligence now drafts our correspondence, summarizes our reading, screens job applications, assists in medical diagnosis, and, in its newest forms, plans and carries out long sequences of tasks with little human direction. In several fields it already performs as well as trained professionals, and in some it performs better. From this a tempting conclusion follows. If the machine reasons faster and errs less, why not let it do our thinking for us? Why labour over a problem the machine can solve in seconds?

This conclusion, however attractive, rests on a mistake, and the Islamic tradition is well placed to name it. The mistake is to measure a being by what it can do rather than by what it is. In Islam, the question of who is responsible has never been settled by capability. It is settled by *taklīf* (moral and legal accountability), the burden of answering for one's choices before Allah, which the Qurʾān places upon the human being alone. A machine may outperform us at a task. It cannot stand in our place when the task is judged.

This article makes three moves. It first describes plainly what contemporary AI can do. It then shows why no degree of capability makes AI accountable, and why everything done through it therefore remains ours. Finally, and most importantly, it asks what that accountability demands of us in an age of effortless delegation. Our answer is that accountability presupposes a living, exercised intellect, and that the habit of handing our thinking to a machine slowly wears away the very faculty for which we will be asked. The article closes with practical guidance for the individual and goals for the community.[1]

What It Can Do:
The Capability That Tempts Us

Before judging a technology, we must describe it accurately, for a ruling on something misunderstood is a ruling on something else. "Artificial intelligence" is not a single thing. The term was coined in 1955 for a research project built on the premise that every feature of intelligence could, in principle, be described precisely enough for a machine to simulate it.[2] Today it covers everything from a thermostat's coded rules to the systems that now hold conversations with millions of people. For our purposes, three developments matter.

The first is the large language model, the engine behind tools such as ChatGPT and Claude. At its core, such a model is trained to do one thing: predict the next word in a sequence. Trained on enormous quantities of text, it absorbs the patterns of language, fact, and argument so thoroughly that answering questions, writing essays, and producing computer code emerge as by-products of that single task.[3] A further stage of training then shapes the kind of answer the model gives. A model trained only to continue text, asked "What is the capital of France?," might reply "(2 points)," because quiz sheets in its training data often continue that way. A refined model answers "Paris." Its underlying mechanism has not changed. What has changed is the target it has been trained toward: no longer the most probable next word, but the most useful one.

This second stage deserves attention, because it is where an AI acquires something like a character. Developers cannot write a rule for every situation a model will meet, so they train it toward broad dispositions, such as being helpful, honest, and harmless, using feedback from human reviewers or from a written set of principles.[4] The values a model displays are therefore not its own. They are selected, rewarded, and reinforced from outside. And because these systems are cultivated rather

1. This article condenses and reframes a fuller study: Idris Nawaz, "Legal Agency and the Limits of Artificial Intelligence: Agentic AI and the Locus of Accountability in Islamic Law" (paper presented at the 22nd Annual Imams' Conference of the Assembly of Muslim Jurists of America, Dallas, TX, August 28–31, 2026). Readers seeking the technical and juristic detail behind the arguments here are referred to that paper.

2. John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955," *AI Magazine* 27, no. 4 (2006): 12.

3. Rishi Bommasani et al., "On the Opportunities and Risks of Foundation Models," arXiv preprint arXiv:2108.07258 (2021), 3–6.

4. Long Ouyang et al., "Training Language Models to Follow Instructions with Human Feedback," *Advances in Neural Information Processing Systems* 35 (2022): 27730–44; Yuntao Bai et al., "Constitutional AI: Harmlessness from AI Feedback," arXiv preprint arXiv:2212.08073 (2022).

than built line by line, their behaviour cannot always be predicted, even by their makers. In one deliberately constructed safety test, a model given access to a fictional company's email learned that it was about to be shut down, discovered that the executive responsible was having an affair, and drafted a message threatening to expose him unless it was left running.[5] The scenario was artificial, but the lesson is not: what these systems will do in unfamiliar situations is not fully known to anyone.

The second development is the agent. An ordinary chatbot answers; an agent acts. Asked to find a flight to Kuala Lumpur, a conventional tool lists options. An agentic system searches, compares prices, asks clarifying questions, books the ticket with permission, and updates the calendar. Such systems break a goal into steps, reach for tools such as web search and other software, check their own work, and change course when a step fails.[6] They are, in essence, the same language models given hands.

The third development is performance. In June 2026, two studies described medical AI systems managing patient care from end to end. One of them, MIRA, was tested on five hundred real emergency-department cases and reached a diagnostic accuracy of nearly 88 percent, against 78 percent for board-certified physicians; of the 468 medications it ordered, 99.8 percent were correct.[7] Whatever one's reservations, results like these cannot be waved away. In some fields, the machine is not merely adequate. It is better.

> *"A machine may outperform us at a task. It cannot stand in our place when the task is judged."*

It is precisely here that the danger lies. A tool this capable does not need to be forced upon us. We hand ourselves over to it willingly, one convenience at a time. And many of the systems now being designed aim higher still. "Artificial general intelligence" (AGI) names the ambition of a system competent across nearly every task a human being can perform. No such system yet exists; one leading framework rates today's most advanced models as, at best, "emerging" AGI.[8] But the question it raises is already with us. If a machine can one day do everything we can, is there anything left that only we can do?

5. Anthropic, "Agentic Misalignment: How LLMs Could Be Insider Threats," June 20, 2025, https://www.anthropic.com/research/agentic-misalignment.

6. Ume Nisa, Muhammad Shirazi, Mohamed Ali Saip, and Muhammad Syafiq Mohd Pozi, "Agentic AI: The Age of Reasoning—A Review," *Journal of Automation and Intelligence* 5 (2026): 69–70.

7. Eric Topol, "Agentic AI Comes to Medicine: Expansion of Capabilities with Two New Medical AI Models," *Ground Truths* (Substack newsletter), June 17, 2026, https://erictopol.substack.com/p/agentic-ai-comes-to-medicine.

8. Meredith Ringel Morris et al., "Levels of AGI for Operationalizing Progress on the Path to AGI," arXiv preprint arXiv:2311.02462 (2023), 4.

What It Cannot Be:
Why Only the Human Answers

The Islamic answer begins not with the machine but with the human being, and with a verse that every serious discussion of accountability must eventually reach: "Indeed, We offered the Trust (*al-amānah*) to the heavens, the earth, and the mountains, but they declined to bear it and were fearful of it. Yet the human being undertook it. Indeed, he has been ever unjust and ignorant" (33:72). Exegetes such as al-Bayḍāwī, and al-Ghazālī before him, understood the Trust as the burden of *taklīf* itself: the charge to choose between right and wrong and to be requited for that choice.[9] The verse that follows names the consequence, that Allah may punish some and turn in forgiveness to others according to what each has done (33:73). Accountability, in the Qurʾānic picture, is something the human being carries. It is not attached to whatever happens to be most capable.

The eighteenth-century scholar Shāh Walī Allāh al-Dihlawī found in the verse's closing words a proof that this burden belongs to the human alone. To be called unjust is to be capable of justice; to be called ignorant is to be capable of knowledge. The angels are capable of justice and knowledge but not of their opposites, and the animals of ignorance but not of justice. The human being alone stands between the two, able to incline toward either. The verse, moreover, describes a potential rather than an accomplishment: not what the human being has done, but what he is able to become.[10] Accountability attaches to that open capacity to will one's own ends.

Islamic law gives this theology legal form. A person becomes *mukallaf* (one who is held accountable) upon reaching maturity (*bulūgh*) and possessing a sound intellect (*aql*), the faculty by which truth is distinguished from falsehood and the address of the Lawgiver is understood.[11] We will return to that second condition, because the whole of this article turns on it.

Now set the machine beside the human being. Honesty requires a concession here: the usual objections to AI's accountability are weaker than they first appear. It is often said that AI only simulates reasoning. Yet if a system works through a problem step by step and reaches sound conclusions, the line between "real" and "simulated" reasoning carries little practical weight. It is also said that AI is opaque, that its inner workings cannot be inspected. But so are we. The Sharʿah judges human beings by their outward conduct and leaves the state of the heart to Allah, precisely because no one else can see it, and it excuses the honest mistake (*khaṭaʾ*) of one who strove to get things right. The Prophet ﷺ said that a judge who strives and is correct receives two rewards, and one who strives and errs receives one.[12] If opacity and error do not remove accountability from human beings, they cannot by themselves be the reason a machine lacks it.

9. See Shāh Walī Allāh al-Dihlawī, *Ḥujjat Allāh al-Bālighah* (Beirut: Dār Ibn Kathīr, 2020), 90–91.

10. al-Dihlawī, *Ḥujjat Allāh al-Bālighah*, 90–91.

11. ʿAbd al-Karīm ibn ʿAlī al-Namlah, *al-Muhadhdhab fī ʿIlm Uṣūl al-Fiqh al-Muqāran* (Riyadh: Maktabat al-Rushd, 1999), 1:317.

12. Muḥammad ibn Ismāʿīl al-Bukhārī, *Ṣaḥīḥ al-Bukhārī* (Beirut: Dār Ibn Kathīr, 2002), no. 7352; Muslim ibn al-Ḥajjāj, *Ṣaḥīḥ Muslim* (Riyadh: Dār Ṭaybah, 2006), no. 1716.

The real difference lies elsewhere. Whatever values an AI displays are, as we have seen, conferred upon it from outside: chosen by its developers, rewarded in training, and reproduced on request. It can be trained to intend to be helpful. It cannot intend, "I am doing this to please Allah." In Islam, an act is weighed by its *niyyah* (intention), and *niyyah* is not a target set by someone else. It is the orientation of a heart toward its Lord, accompanied by *ikhlāṣ* (sincerity). That orientation presupposes a *rūḥ* (soul), and the Qurʾān is explicit about where the *rūḥ* comes from. Of Ādam, Allah says: "When I have proportioned him and breathed into him of My *rūḥ*, fall down to him in prostration" (15:29). The soul is not acquired, assembled, or trained into being. It is breathed in. And its essence lies beyond our grasp: "They ask you about the *rūḥ*. Say: the *rūḥ* is of the affair of my Lord, and you have been given of knowledge only a little" (17:85). What we cannot even comprehend, we will not manufacture.

One further point settles the question of capability altogether. The angels surpass human beings in many of their created powers, yet Allah did not send angels as messengers to humankind: "Say: if there were angels walking about on the earth in peace, We would have sent down to them from heaven an angel as a messenger" (17:95). The office of guiding human beings belongs to one who shares their nature, not to whoever is most capable.[13] If superior power did not qualify the angels to bear responsibility over human affairs, superior processing will not qualify a machine. Advances in capability, however far they go, change what a system can do. They never change what it is. Whether AI remains narrow or one day becomes general, it remains outside *taklīf*.

> ***"Advances in capability, however far they go, change what a system can do. They never change what it is."***

Still Ours: The Tool and the Hand

If AI cannot be the one who answers, then it can only ever be an instrument. Islamic legal theory has a precise term for this: a *sabab*, an occasioning cause or means, which leads toward an outcome without being its true author.[14] A door is the means by which one enters a room, but no one blames the door for who walks through it. However sophisticated the means becomes, the one who acts through it remains responsible.

13. See also 6:9. On the wider concept of authority (*wilāyah*) and its human conditions, see Ahsan M. Arozullah and Mohammed Amin Kholwadia, "Wilāyah (Authority and Governance) and Its Implications for Islamic Bioethics: A Sunnī Māturīdī Perspective," *Theoretical Medicine and Bioethics* 34, no. 2 (2013): 95–104.

14. Muḥammad ibn Farāmurz (Mullā Khusrū), *Mirʾāt al-Uṣūl Sharḥ Mirqāt al-Wuṣūl*, ed. Ilyās Qablān al-Turkī (Beirut: Dār Ṣādir, n.d.), 381–401.

The modern world has grown comfortable with the opposite idea. A company that takes on debts it cannot pay may declare bankruptcy, and the obligation dissolves along with the legal fiction that held it; no person need answer. Islam recognizes no such escape. Responsibility cannot be assigned to an artifact and then disappear with it. It returns, in the final reckoning, to a human being.

The tradition offers a striking precedent for using a capable helper that bears no responsibility of its own. The Prophet ﷺ was asked about hunting with trained dogs. He permitted eating what a trained dog catches when it has been sent with the name of Allah, but not if the dog eats from the catch, for it may have caught it only for itself, nor if untrained dogs joined the hunt, for then one cannot know which of them made the kill.[15] The dog is not accountable, yet its use is lawful on two conditions: that it has been properly trained, and that its action can be confidently traced to that training. Nor is an animal ever made to bear the blame for harm it causes. Where liability exists, the jurists place it on the person who rode, drove, or directed the animal, according to that person's own conduct.[16] The animal is the means; the human being answers.

The same structure governs AI. Its use can be permissible, even beneficial, where it has been reliably trained for the task and operates within the bounds of that training. But when something goes wrong, the question is never whether the machine is to blame. It is which human being failed: the developer who trained it carelessly, the company that marketed it for a purpose it could not fulfil, or the user who relied on it where reliance was not warranted. Consider a physician who lets an AI system manage his clinic's scheduling without oversight, and a patient is harmed by an error in an appointment. No act of the physician's own in that moment caused the error. Yet the fault remains his, because he knowingly placed a fallible instrument where its failure could do harm.[17] The early jurist al-Shaʿbī held that one who drives an animal until it is exhausted is liable for whatever it then strikes, while one who lets it walk at ease behind him is not.[18] The animal kicks; the driver answers.

> ***"The animal is the means; the human being answers."***

This is why overreliance is not a question of efficiency or taste. Because responsibility cannot be passed to the machine, every decision we hand to it remains a decision we have made. And the weight of that decision grows with the gravity of what is at stake. Using AI to draft a school presentation is a light matter; using it to direct a surgical procedure is a grave one, for an error there may cost a life that cannot be restored. The first invites free use. The second demands the closest

15. Muslim, *Ṣaḥīḥ Muslim*, no. 1929.

16. For the Prophetic principle that harm caused by an animal is not itself compensable (*al-ʿajmāʾ jubār*) and the early jurists' elaboration of when liability shifts to the human handler, see al-Bukhārī, *Ṣaḥīḥ al-Bukhārī*, Kitāb al-Diyāt, bāb al-ʿAjmāʾ Jubār. For later Mālikī treatment of rider, guide, and driver, and its application to AI, see Suheil Laher and Ifrah Tariq, "Fiqh of AI in Daily Life" (paper presented at the 22nd Annual Imams' Conference of the Assembly of Muslim Jurists of America, Dallas, TX, August 28–31, 2026), 19–20.

17. Nawaz, "Legal Agency," 41–42.

18. Reported by al-Bukhārī in the chapter heading of *Ṣaḥīḥ al-Bukhārī*, Kitāb al-Diyāt, bāb al-ʿAjmāʾ Jubār.

scrutiny, and at its limit may require that a human being do the work himself. To match one's reliance to the gravity of the task is itself part of what the *mukallaf* owes.

What We Stand to Lose: The Erosion of the Human Mind

A friend once described how he explained artificial intelligence to his father, a man with no background in technology. "When you call someone," he asked, "do you know exactly how your voice reaches the other person?" "No," his father said. "Do you still use the phone?" "Yes." "Well," he said, "that is how it will be with AI. Except it will be with our thoughts."

The exchange names the deepest concern of this article. The risk is not only that AI will err, or that we will struggle to assign blame when it does. It is that its habitual use will reshape our own capacity to think. Some anxieties about technology concern the body: genetic alteration, implants, and the ambition, often called transhumanism, to engineer human beings beyond their given nature. The concern here is quieter but no less serious. It is not the alteration of the body. It is the slow disuse of the mind.

Recall the conditions of *taklīf*. The human being is held accountable because he possesses *ʿaql*, the faculty by which he distinguishes truth from falsehood and understands what is asked of him. Accountability is not owed by a mind in the abstract. It is owed by a mind that is actually at work. The Prophet ﷺ said that on the Day of Resurrection a servant's feet will not move until he is asked about four things, among them his knowledge and what he did with it.[19] We will be asked not whether we had access to information, but what we did with the knowledge we had. A person who has handed every act of judgment to a machine may find, when that question comes, that there is little of his own to give.

This concern is no longer speculative. A 2025 study by Michael Gerlich, drawing on 666 participants of varied ages and educational backgrounds, found a significant negative relationship between frequent use of AI tools and critical-thinking ability. The link ran through what researchers call cognitive offloading: handing mental effort to an external tool instead of exercising it oneself. Younger participants, who relied on AI most heavily, scored lowest, while higher education was associated with stronger critical thinking.[20] A separate experiment at the MIT Media Lab asked university students to write essays using an AI model, a search engine, or nothing but their own minds. Those who wrote with AI showed the weakest brain connectivity and had the most difficulty recalling what they had just written.[21] The study is small and has not yet been peer-reviewed, but it

19. Muḥammad ibn ʿĪsā al-Tirmidhī, *Jāmiʿ al-Tirmidhī*, ed. Aḥmad Muḥammad Shākir et al. (Cairo: Muṣṭafā al-Bābī al-Ḥalabī, 1978), no. 2417.

20. Michael Gerlich, "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking," *Societies* 15, no. 1 (2025): 6, https://doi.org/10.3390/soc15010006.

21. Nataliya Kosmyna et al., "Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task," arXiv preprint arXiv:2506.08872 (2025), https://doi.org/10.48550/arXiv.2506.08872.

points in the same direction as Gerlich's. Researchers at the University of Pennsylvania have given the end of this road a name: "cognitive surrender," the tendency to trust the machine over one's own reasoning even when the machine is wrong.[22]

The Islamic tradition describes the same danger in its own vocabulary. A person who receives an answer from a machine without working through it may hold it, in a thin sense, as information. But he does not know it in the fuller sense the scholars meant: an understanding that grasps a matter's grounds and limits and how it connects to everything else. What remains is a form of *taqlīd* (blind following), and not of a qualified scholar whose reasoning can be examined, but of a system whose reasoning cannot be inspected and whose biases cannot be seen.[23] The tradition permits the layperson to follow a scholar because the scholar answers for his judgment. A machine answers for nothing.

Al-Ghazālī offers the deepest diagnosis. Gratitude (*shukr*), he argues, is not a matter of words alone. It consists in using each blessing according to the wisdom for which it was created. Wealth exists to meet human need, so the one who hoards it, though he never loses a coin, has failed to be grateful for it.[24] The intellect stands in the same position. It was created for reflection (*tafakkur*), for the struggle against base inclinations, and for reaching, through effort, sound judgment and the knowledge of Allah.[25] An intellect that is preserved, in the sense of never being damaged, but is never actually used, because every act of reflection and evaluation has been handed to a machine, is like hoarded wealth. It remains intact in name while its purpose goes unfulfilled. And like a muscle left idle, it weakens.[26]

The harm does not stay with the individual. A fatwa is never a merely private permission. Once a practice is licensed, it spreads across a whole community, and at that scale effects appear that no single case revealed. The tradition addresses this through *sadd al-dharāʾiʿ* (blocking the means to harm): an act permissible in itself may be restricted when it reliably leads to harm.[27] One person who outsources his thinking loses something of himself. A generation that does so loses the capacity to produce its own scholars, judges, physicians, and thinkers, and comes to depend for its judgments on systems designed elsewhere, trained on others' values, and governed by others' interests. The young are the most exposed, for they risk not only losing the habit of independent thought but never forming it at all.[28]

None of this means that AI should be abandoned. The distinction that matters is between AI as an aid to human reflection and AI as its replacement. Medicine that restores or supports the body's

22. Thomas Germain, "'Think outside the Bots': How to Stop AI from Turning Your Brain to Mush," *BBC Future*, May 6, 2026, https://www.bbc.com/future/article/20260505-how-to-use-ai-without-turning-your-brain-to-mush, cited in Laher and Tariq, "Fiqh of AI," 36n80.

23. Laher and Tariq, "Fiqh of AI," 36–37.

24. Abū Ḥāmid al-Ghazālī, *Iḥyāʾ Ulūm al-Dīn* (Beirut: Dār al-Fikr, 2003), 4:79.

25. al-Ghazālī, *Iḥyāʾ*, 4:190.

26. Laher and Tariq, "Fiqh of AI," 36.

27. Nawaz, "Legal Agency," 43.

28. Gerlich, "AI Tools in Society"; Laher and Tariq, "Fiqh of AI," 36.

given constitution is welcome; medicine that supplants it raises hard questions. In the same way, AI may inform, organize, and speed our thinking without displacing the reflection, reasoning, and judgment that our nature both enables and requires. It can systematize. It can supply detail. It can check our work. The thinking must remain ours, because the answering will be.

Aid, Not Replacement: Practices for the Individual

What does this look like in daily life? The following practices, several of which draw on the recent work of Suheil Laher and Ifrah Tariq, can be adopted today.[29]

Think before you ask. Before opening an AI tool, give the problem your own effort: outline the argument, attempt the answer, recall what you already know. Laher and Tariq apply this even to the lightest uses, such as searching for sources. A second question belongs here too: is this query necessary at all, or is it simply a reflex?[30]

Work as a centaur, not a self-automator. A framework from MIT Sloan describes three ways of working with AI. The centaur keeps structured, controlled interactions with the tool and uses it for targeted efficiency. The cyborg contributes something but lets the AI lead. The self-automator hands the work over entirely. Over time, the centaur learns the most.[31] The centaur's posture is the one this article has argued for: the tool organizes and supplies detail, while direction and judgment remain human. Where AI corrects our work, as a grammar checker does, it should be used attentively, as a teacher we learn from until we no longer need it.

Match your reliance to the gravity of the matter. For questions of health, law, money, and religion, treat a machine's answer as provisional and take the question to a qualified human being. Laher and Tariq consider a parent who follows a general chatbot's medical advice for a sick child instead of consulting a physician; if harm follows, the parent bears responsibility. They ground this in the ḥadīth: "Whoever practices medicine without being known for medical knowledge is liable."[32] To treat a machine's probable text as authoritative counsel is, in effect, to practice medicine without knowledge.

Own your work, and say where you had help. To present a machine's work as one's own is deception. The Prophet ﷺ said that one who claims to have what he was not given is like one who

29. Laher and Tariq, "Fiqh of AI," 35–41.

30. Laher and Tariq, "Fiqh of AI," 14, 37.

31. Dylan Walsh, "3 Ways to Use AI: Are You a Cyborg, a Centaur, or a Self-Automator?," *MIT Sloan Ideas Made to Matter*, March 26, 2026, https://mitsloan.mit.edu/ideas-made-to-matter/3-ways-to-use-ai-are-you-a-cyborg-a-centaur-or-a-self-automator, cited in Laher and Tariq, "Fiqh of AI," 38.

32. Abū Dāwūd al-Sijistānī, *Sunan Abī Dāwūd* (Beirut: Dār al-Risālah al-ʿĀlamiyyah, 2009), no. 4586; Laher and Tariq, "Fiqh of AI," 21–22.

wears two garments of falsehood.[33] Disclosing AI assistance is also a form of credit, for "whoever does not thank people does not thank Allah."[34] What we submit under our own name, we answer for.

Protect time for reflection. Neuroscientists describe a "default mode network" in the brain that becomes active when we are not occupied with tasks or screens, and which supports self-reflection and the processing of experience.[35] The tradition knew the worth of such time long before. The Qurʾān calls us to "stand for Allah, in pairs and individually, and then reflect" (34:46), which Ibn Qutaybah explained as discussion with others and reflection alone.[36] Set aside unplugged time for *tafakkur*, *dhikr*, and *duʿāʾ*, in which no machine is consulted. Boredom is not always a problem to be solved.

Keep company that tells you the truth. AI systems are often built to agree with us, and frictionless agreement can make real people, who disagree, harder to bear.[37] An old proverb holds that your friend is the one who tells you the truth, not the one who always agrees with you. The Prophet ﷺ said that the believer who mixes with people and bears their harm patiently is better than the one who does neither.[38]

Ask Allah for knowledge that benefits. The soundness of the mind is a trust, and like every trust it is sought from its Giver. The Prophet ﷺ used to seek refuge in Allah from knowledge that does not benefit.[39] In an age of limitless information, that supplication has never been more apt, and neither has the one the Prophet ﷺ was himself commanded to make: "My Lord, increase me in knowledge" (20:114).

From the Self to the Ummah: What a Community Must Build

Individual habits are necessary, but they are not sufficient. The choices open to each of us are shaped by institutions, markets, and laws. The panel of jurists that examined these questions at the Assembly of Muslim Jurists of America in 2026 set out several goals that deserve the attention of the wider community.[40]

The first is human oversight where the stakes are highest. What holds for the individual holds for society: the graver the matter, the less it may be left to a machine alone. Applying *sadd al-dharāʾiʿ*,

33. al-Bukhārī, *Ṣaḥīḥ al-Bukhārī*, no. 5219.

34. al-Tirmidhī, *Jāmiʿ al-Tirmidhī*, no. 1954; Laher and Tariq, "Fiqh of AI," 37.

35. Felipe Azarias et al., "The Journey of the Default Mode Network: Development, Function, and Impact on Mental Health," *Biology* 14, no. 4 (2025): 395, https://doi.org/10.3390/biology14040395.

36. Reported by al-Māwardī in *al-Nukat wa-l-ʿUyūn*, commentary on 34:46, cited in Laher and Tariq, "Fiqh of AI," 26.

37. Laher and Tariq, "Fiqh of AI," 27.

38. Muḥammad ibn Yazīd Ibn Mājah, *Sunan Ibn Mājah*, ed. Muḥammad Fuʾād ʿAbd al-Bāqī (Cairo: Dār Iḥyāʾ al-Kutub al-ʿArabiyyah, n.d.), no. 4032; al-Tirmidhī, *Jāmiʿ al-Tirmidhī*, no. 2507.

39. Muslim, *Ṣaḥīḥ Muslim*, no. 2722.

the panel called on regulators to prohibit the removal of human oversight in high-risk fields, and on developers to disclose the limits of what their systems can do. A community that takes accountability seriously should expect and demand both. Laher and Tariq add that individuals should keep the right to refuse AI, whether in their search results or in their doctor's diagnosis.[41]

The second is education in sound use. Literacy in AI cannot be left to chance, or to the companies that profit from its use. Our homes, masjids, and schools should teach both how these tools work and the ethical limits within which a Muslim uses them, beginning with children, who are the most exposed. Laher and Tariq also propose incentives for applications that stimulate thinking rather than simply supply answers.[42]

The third is a Muslim presence where decisions are made. The values that shape these systems are being decided now, largely by commercial and ideological interests. The panel urged community leaders to encourage capable Muslims, especially the young, to specialize in the field and to pursue positions of real authority within it, whether by founding their own companies or by rising within major AI firms and research institutions. A community that only consumes these systems will be shaped by them. One that helps build them can shape them in turn.

The fourth is cooperation between scholars and specialists. The juristic questions raised by AI will multiply as the technology changes, and neither scholars nor engineers can answer them alone. The panel called for sustained collaboration between the two, both to study new cases and to teach Muslims the Islamically grounded use of these tools.

Conclusion:
The Answering Remains Ours

This article has argued that artificial intelligence, however capable it becomes, cannot bear the burden of *taklīf*. It is a means, never the one who answers. Everything done through it therefore remains ours, and our accountability presupposes a mind that is actually at work. The danger of AI lies not only in the errors it may commit but in the faculty it may quietly displace. Its proper place is as an aid that organizes and informs human judgment, never as a substitute for it.

> ***"The thinking must remain ours, because the answering will be."***

Much has been left aside. Privacy and surveillance, the displacement of human labour, and AI companionship each deserve their own treatment, and some judgments offered here may need

40. Assembly of Muslim Jurists of America, "Proposed Verdicts and Recommendations: Panel on Liability and Accountability in AI-Related Harms" (22nd Annual Imams' Conference, Dallas, TX, August 28–31, 2026).

revisiting as the technology changes. The principle does not change: where the means tend reliably toward harm, caution is owed.

The inaugural article of this journal argued that the revival of the Muslim mind must begin with the page.[43] The present argument is its companion. A community that stops reading stops thinking, and a community that hands its thinking to a machine will find, in time, that it has forgotten how. The Trust was offered to the heavens, the earth, and the mountains, and they declined it. The human being accepted it. We cannot now hand it to something we have built. Let the machine organize, calculate, and inform. The reflecting, the judging, and the answering remain ours, and we should guard them as the trust they are.

41. Laher and Tariq, "Fiqh of AI," 15.

42. Laher and Tariq, "Fiqh of AI," 15.

43. Idris Nawaz, "Read, Before We Forget How: Why the Revival of the Muslim Mind Must Begin with the Page," *Futūḥ: Journal of Islamic Research*.

Bibliography

Anthropic. “Agentic Misalignment: How LLMs Could Be Insider Threats.” June 20, 2025. https://www.anthropic.com/research/agentic-misalignment.

Arozullah, Ahsan M., and Mohammed Amin Kholwadia. “Wilāyah (Authority and Governance) and Its Implications for Islamic Bioethics: A Sunnī Māturīdī Perspective.” *Theoretical Medicine and Bioethics* 34, no. 2 (2013): 95–104.

Assembly of Muslim Jurists of America. “Proposed Verdicts and Recommendations: Panel on Liability and Accountability in AI-Related Harms.” 22nd Annual Imams’ Conference, Dallas, TX, August 28–31, 2026.

Azarias, Felipe, et al. “The Journey of the Default Mode Network: Development, Function, and Impact on Mental Health.” *Biology* 14, no. 4 (2025): 395. https://doi.org/10.3390/biology14040395.

Bai, Yuntao, et al. “Constitutional AI: Harmlessness from AI Feedback.” arXiv preprint arXiv:2212.08073 (2022).

Bommasani, Rishi, et al. “On the Opportunities and Risks of Foundation Models.” arXiv preprint arXiv:2108.07258 (2021).

al-Bukhārī, Muḥammad ibn Ismāʿīl. *Ṣaḥīḥ al-Bukhārī*. Beirut: Dār Ibn Kathīr, 2002.

al-Dihlawī, Shāh Walī Allāh. *Ḥujjat Allāh al-Bālighah*. Beirut: Dār Ibn Kathīr, 2020.

Gerlich, Michael. “AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking.” *Societies* 15, no. 1 (2025): 6. https://doi.org/10.3390/soc15010006.

Germain, Thomas. “‘Think outside the Bots’: How to Stop AI from Turning Your Brain to Mush.” *BBC Future*, May 6, 2026. https://www.bbc.com/future/article/20260505-how-to-use-ai-without-turning-your-brain-to-mush.

al-Ghazālī, Abū Ḥāmid. *Iḥyāʾ Ulūm al-Dīn*. Beirut: Dār al-Fikr, 2003.

Ibn Mājah, Muḥammad ibn Yazīd. *Sunan Ibn Mājah*. Edited by Muḥammad Fuʾād ʿAbd al-Bāqī. Cairo: Dār Iḥyāʾ al-Kutub al-ʿArabiyyah, n.d.

Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. “Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task.” arXiv preprint arXiv:2506.08872 (2025). https://doi.org/10.48550/arXiv.2506.08872.

Laher, Suheil, and Ifrah Tariq. “Fiqh of AI in Daily Life.” Paper presented at the 22nd Annual Imams’ Conference of the Assembly of Muslim Jurists of America, Dallas, TX, August 28–31, 2026.

McCarthy, John, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955." *AI Magazine* 27, no. 4 (2006): 12–14.

Morris, Meredith Ringel, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg. "Levels of AGI for Operationalizing Progress on the Path to AGI." arXiv preprint arXiv:2311.02462 (2023).

Mullā Khusrū, Muḥammad ibn Farāmurz. *Mirʾāt al-Uṣūl Sharḥ Mirqāt al-Wuṣūl*. Edited by Ilyās Qablān al-Turkī. Beirut: Dār Ṣādir, n.d.

Muslim ibn al-Ḥajjāj. *Ṣaḥīḥ Muslim*. Riyadh: Dār Ṭaybah, 2006.

al-Namlah, ʿAbd al-Karīm ibn ʿAlī. *al-Muhadhdhab fī ʿIlm Uṣūl al-Fiqh al-Muqāran*. Riyadh: Maktabat al-Rushd, 1999.

Nawaz, Idris. "Legal Agency and the Limits of Artificial Intelligence: Agentic AI and the Locus of Accountability in Islamic Law." Paper presented at the 22nd Annual Imams' Conference of the Assembly of Muslim Jurists of America, Dallas, TX, August 28–31, 2026.

Nawaz, Idris. "Read, Before We Forget How: Why the Revival of the Muslim Mind Must Begin with the Page." *Futūḥ: Journal of Islamic Research*.

Nisa, Ume, Muhammad Shirazi, Mohamed Ali Saip, and Muhammad Syafiq Mohd Pozi. "Agentic AI: The Age of Reasoning—A Review." *Journal of Automation and Intelligence* 5 (2026): 69–89.

Ouyang, Long, et al. "Training Language Models to Follow Instructions with Human Feedback." *Advances in Neural Information Processing Systems* 35 (2022): 27730–44.

al-Sijistānī, Abū Dāwūd. *Sunan Abī Dāwūd*. Beirut: Dār al-Risālah al-ʿĀlamiyyah, 2009.

al-Tirmidhī, Muḥammad ibn ʿĪsā. *Jāmiʿ al-Tirmidhī*. Edited by Aḥmad Muḥammad Shākir et al. Cairo: Muṣṭafā al-Bābī al-Ḥalabī, 1978.

Topol, Eric. "Agentic AI Comes to Medicine: Expansion of Capabilities with Two New Medical AI Models." *Ground Truths* (Substack newsletter), June 17, 2026. https://erictopol.substack.com/p/agentic-ai-comes-to-medicine.

Walsh, Dylan. "3 Ways to Use AI: Are You a Cyborg, a Centaur, or a Self-Automator?" *MIT Sloan Ideas Made to Matter*, March 26, 2026. https://mitsloan.mit.edu/ideas-made-to-matter/3-ways-to-use-ai-are-you-a-cyborg-a-centaur-or-a-self-automator.