



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

Fiqh of AI in Daily Life

Dr. Suheil Laher
Dr. Ifrah Tariq

Dr. Suheil Laher has an MA in Religious Studies from Boston University and a PhD in Arabic and Islamic Studies from Harvard University. He previously served as Academic Dean at Fawakih Institute where he made significant contributions to the classical Arabic curriculum. He has previously taught at Harvard, Brandeis and Boston Universities, and has also taught courses at the Ella Collins Institute and the Bukhari Institute. He has benefited from a number of teachers of traditional Islamic disciplines, in the U.S. and abroad, and has ijazahs in Islamic theology ('aqidah), Islamic law (fiqh), hadith and other subjects.

He conducts research in Islamic theology, law, Quran and hadith studies. He has published academic articles, book chapters, encyclopedia articles, and scholarly translations, and has presented at academic conferences in the US and abroad. He served as Muslim Chaplain at MIT for over ten years, an experience that helped lead him to believe strongly in the value of both inter-religious and intra-Muslim dialogue, with a view to better understanding and fruitful cooperation to help make God's earth a better place.

Dr. Ifrah Tariq is a Software Engineer at Google, where she conducts research on foundational Large Language Models (LLMs). She holds a PhD in Computational and Systems Biology from MIT; her doctoral thesis focused on "Biologically Interpretable Representation Learning for Mechanistic Insights into Cancer Immunotherapy Resistance." Ifrah also earned a Master's degree in Computational Biology from Cambridge University and holds an undergraduate degree in Molecular Biology and Biochemistry. Currently, her professional and personal research focuses on advancing generalizable models and AI at scale, with a dedicated emphasis on ethical AI and its utility in medicine and daily life.

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

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




TABLE OF CONTENTS

1. Introduction to Artificial Intelligence	4
1.1 Historical Background	4
1.2 Types of AI	6
1.3 Theological and Ethical Dimensions	7
2. Foundational Fiqh Maxims.....	10
A) The Five Major Fiqh Maxims.....	10
I. Intentions and Goals <i>الأمر بمقاصدها</i>	10
II. Certainty is not displaced by Doubt <i>اليقين لا يزول بالشك</i>	11
III. Hardship brings Leniency <i>المشقة تجلب التيسير</i>	11
IV. No Harming, nor Reciprocation of Harm <i>لا ضرر ولا ضرار</i>	11
V. Customary Norms are Arbiters <i>العادة محكمة</i>	12
B) The Higher Objectives	12
I. Religion (<i>dīn</i>)	13
II. Life (<i>nafs</i>)	13
III. Intellect (<i>'aql</i>)	13
IV. Family and lineage (<i>nasl</i>)	13
V. Property (<i>māl</i>)	13
C) Maxims Specific to AI Regulation	14
Specific Issues	17
Liability and Accountability	17
A) Technical Background	17
B) <i>Fiqhī</i> Dimensions	19
AI Companions and Virtual Relationships	24
A) Technical Background	24
B) <i>Fiqhī</i> Dimensions.....	26
Deepfakes and Digital Manipulation	32
A) Technical Background	32
B) <i>Fiqhī</i> Dimensions.....	33
Daily Conduct and Moral Harm	35
A) Technical Background	35
B) <i>Fiqhī</i> Dimensions	36
Conclusion	44

الحمد لله والصلاة والسلام على رسول الله

The world has changed drastically within the last few years due to an explosion of new Artificial Intelligence (AI) technologies and related phenomena including chatbots, AI language translation, deepfakes, AI-assisted warfare, and addictive social media algorithms. How can Muslims interact with these technologies in an ethical way that does not compromise our beliefs or values? This paper is an initial attempt to chart these waters by providing some general guidelines as well as engaging specific current issues. We start with an introduction to the history of AI and its current forms, and some related theological and ethical issues. This is followed by a section on how fiqh maxims and higher objectives of sharia interface with AI. Finally, we turn to specific examples across four broad topics.

Introduction to Artificial Intelligence

1.1 Historical Background

Understanding the evolution of artificial intelligence is essential to contextualize its current form. Artificial intelligence is a broad field encompassing techniques that enable machines to perform tasks or exhibit behaviors typically associated with human intelligence. In the 1950s, Alan Turing proposed a foundational perspective on machine intelligence, suggesting that a machine could be considered intelligent if its behavior was indistinguishable from that of a human. He established the Turing Test, in which a human evaluator engaged in a text-based conversation with two unseen participants: one human and one machine. If, after the interaction, the evaluator could not reliably distinguish between the machine and the human, the machine would “pass” the test.

The key principle of the Turing Test is not to define a machine’s capacity to think or understand as humans do, but rather to assess whether its responses are sufficiently convincing to appear human-like. This fundamentally shifted the focus from internal processes (the “black box” underlying the machine’s decisions) toward observable behavior (its output).

This test highlighted the philosophical debate shift from “can machines think” to something that could be more directly interrogated: “can machines outperform humans in a specified task”?¹

Early AI development focused on Symbolic AI, which relied on logical rules and definitions to represent knowledge and ‘understanding’. This was prominent in expert systems in the 70s and 80s, which were used to diagnose diseases (e.g., MYCIN, INTERNIST-I) and IBM’s Deep Blue, a supercomputer that relied on human-coded heuristics designed to play chess.

Early AI development focused on Symbolic AI, which utilized explicit logical rules and human-coded heuristics to represent knowledge. This paradigm dominated the 1970s and 1980s, exemplified by expert systems such as MYCIN and INTERNIST-I for medical

¹The Turing-centric approach is not without its critics. It presupposes that natural-language proficiency is a sufficient indicator of human-level cognition, a view that prioritizes behavioral mimicry over semantic understanding.. Philosopher John Searle articulated this in the 1980s in his “Chinese Room” thought experiment. For more details, see:

<https://doi.org/10.1017/S0140525X00005756>

diagnosis, and culminating in IBM's Deep Blue, which achieved mastery in chess through rule-based search algorithms.

As computational power increased, the field shifted toward mathematical and statistical approaches. Rather than relying on manually defined rules, Machine Learning (ML) enables algorithms to identify patterns within large datasets. Supervised learning models, for instance, are trained to minimize predictive error against a human-labeled "ground truth." However, this approach faced two critical bottlenecks: a heavy dependence on high-quality labeled data and limited generalizability across disparate domains.

To overcome these constraints, the field evolved toward self-supervised learning. In this paradigm, models generate their own supervisory signals from raw, unlabeled data—for example, by predicting a masked word based on its surrounding context. This process necessitates that models develop deep, abstract representations of the data's underlying structure. This shift was facilitated by encoder-decoder architectures, where an encoder compresses raw input (e.g. text strings, numbers, etc.) into a latent, fixed-length vector—a data structure that captures a distilled abstraction—before a decoder generates an output for a specific downstream task. Mirroring the philosophy of the Turing Test, the efficacy of these representations is validated by downstream performance rather than the interpretability of the internal weights.

A pivotal shift occurred in 2017 with the introduction of the Transformer architecture.² Central to this innovation is the "attention" mechanism, which allows models to compute the relative importance of every input feature in relation to every other feature, rather than treating inputs as independent variables. By utilizing relational matrices, Transformers move beyond simple feature summation, allowing for more nuanced contextual modeling. This shift has yielded significant empirical improvements in pattern recognition across diverse domains.

The current trajectory of AI is defined by the rise of 'foundational models'—large-scale systems trained on vast, diverse datasets that leverage transfer learning to adapt to a wide array of tasks. This is an active and burgeoning area of research; these models are already demonstrating remarkable utility, offering unprecedented efficiency in synthesizing complex information and automating sophisticated workflows. However, their deployment is not without significant operational challenges. Many of these models suffer from performance degradation due to data distribution shifts, as they are typically trained in an 'offline' capacity—relying on a static, historical snapshot of information that may not account for real-time environmental changes. While 'online' learning, which allows for continuous adaptation, is a promising path forward, it introduces substantial risks regarding computational cost and system stability.

Crucially, the inherent power of these tools does not diminish the need for human expertise. Because foundational models are designed to be highly articulate, scholars and practitioners often overestimate their underlying reasoning capabilities, leading to an over-reliance that can result in critical errors. Consequently, the utility of these systems is heavily dependent on the presence of domain experts who can critically evaluate and interpret their outputs. Furthermore, there is often limited corporate incentive to bridge this 'perception-performance' gap; organizations frequently prioritize widespread adoption over

² Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention is all you need." *Advances in neural information processing systems* 30 (2017). Searle, John R. "Minds, brains, and programs." *Behavioral and brain sciences* 3, no. 3 (1980): 417-424; <https://doi.org/10.1017/S0140525X00005756>

transparent capability assessment, offering only superficial disclaimers regarding the possibility of errors. This is complicated by the fact that defining the boundaries of these models is an immense technical challenge. While researchers can measure performance on known benchmarks, it is significantly harder to quantify what we do not yet understand about a model's failure modes. As this remains an active and evolving field of research, there are no simple 'safety' metrics that guarantee reliability. Consequently, the burden of discretion falls upon the user; until these systems are more predictable, all parties must exercise professional caution, treating model outputs as provisional insights that require rigorous expert validation rather than as authoritative truths.

1.2 Types of AI

AI systems can be broadly characterized based on their technical objective, model architecture, and operational role.

A. By Learning Objective:

- a. Supervised Learning: Systems trained on labeled "ground truth" datasets to perform classification or regression tasks (e.g. medical diagnostics).
- b. Unsupervised Learning: Systems designed to discover latent structures, clusters, or distributions within raw data without human-provided labels. (e.g. models to cluster different kinds of customers)
- c. Self-Supervised Learning: Systems that derive patterns from raw, unlabeled data by predicting hidden or future segments. This is a primary driver behind modern foundational models. (e.g. LLMs)
- d. Reinforcement Learning: Systems that learn through iterative feedback, optimizing for a reward signal. This is essential for aligning models with human intent during fine-tuning. (e.g. AlphaGo - an algorithm for playing the game Go)

B. By Development Objective:

- a. Foundational Models: Large-scale, general-purpose neural networks pre-trained on expansive datasets to serve as the shared infrastructure for a wide range of tasks, effectively acting as a starting point rather than a purpose-built tool. (e.g. GPT)
- b. Specialized Models: Targeted systems designed for a specific domain or objective, created either by training a model from scratch—building its knowledge base from a blank slate—or by fine-tuning a foundational model on curated, domain-specific data. (e.g. AlphaFold - an algorithm for predicting protein structure based on amino acid sequences)

C. By Functional Role:

- a. Predictive Systems: Tools that categorize or forecast. These systems use supervised classification to flag content—such as hate speech or policy violations—based on historical patterns. While efficient for large-scale filtering, they often require human oversight for ambiguous cases, such as distinguishing between irony and literal threats. (e.g. content management)
- b. Generative Systems: Tools that synthesize new content. These function as conversational interfaces, leveraging the model's probabilistic "understanding" of language to generate responsive text. Their utility relies on the model's generalist knowledge, though they require expert validation to mitigate

- hallucinations. (e.g. chatbots)
- c. Agentic Systems: Tools designed for autonomous task execution. These systems use a reasoning engine to plan, interact with software tools (like APIs or browsers), and iterate until a goal is met. In this context, engineers use the term reasoning to describe the system's capacity to decompose a high-level objective into a logical sequence of sub-tasks and dynamically adjust its approach based on the success or failure of previous steps. (e.g. automated research agents)

1.3 Theological and Ethical Dimensions

Before exploring the fiqh of AI, it is useful to take a step back to examine the term itself, in order to have a clearer conception of it and related issues. Although, as the ulema say, "Definitions of technical terms cannot be contested" (لا مشاحة في الاصطلاح), this principle is subject to the stipulation that there is no contravention of Islamic theology or principles. Names and terms can have theological significance. For example, the Prophet ﷺ:

- condemned humans using the title "King of Kings."³
 - predicted that some people would drink wine while calling it something else.⁴
- Another hadith intimates that some words have a detrimental psychological effect:

لا يَقُولَنَّ أَحَدُكُمْ: خَبِثْتُ نَفْسِي، وَلَكِنْ لِيَقُلْ: لَقِستَ نَفْسِي.

"None of you should say, 'I am bad,' but rather say, 'I am discontented.'"⁵

For the matter at hand, it is worth noting that a number of leading figures in the field are atheists, and some of them pursue AI as a means to rival God.

"Who is more of a wrongdoer than one who tries to create like My creation? Let them then create a seed, or let them create a tiny ant!"⁶

"The people with the most severe punishment on the Day of Resurrection will be those who seek to imitate the creation of Allah."⁷

For example, American computer scientist Ray Kurzweil (b. 1948) predicts (fantasizes) that humans would become immortal by uploading the contents of their brain to computers. Marvin Minsky (d. 2016), one of the fathers of AI, described the human brain as a mere "meat machine," and believed that death would be overcome through science. Although he himself rejected a comparison between humans and God (despite being an atheist), he saw the AI enterprise as similar to the role of Prometheus, a figure from Greek mythology who, by bringing fire to humans from 'the gods' gave humans access to "mysteries that were once reserved only for the gods."⁸ In a similar vein, when Jeff Bezos launched his AI company in 2025, he named it "Project Prometheus." In some circles, AI is being portrayed as quasi-divine, and it is true that, "the very notion of creating artificial consciousness, or a digital "soul," taps into fundamental theological questions about the

³ Bukharī.

⁴ Abū Dāwūd, Ibn Mājah.

⁵ Bukharī.

⁶ Bukharī, Muslim.

⁷ Bukharī, Muslim.

⁸ Jeffrey Mishlove's interview with Minsky, <https://www.organism.earth/library/document/mind-as-society>, accessed 5/30/2026.

nature of life, existence, and the relationship between creator and creation.”⁹ TESCREAList is an acronym¹⁰ used to describe those who think AI will become an all-powerful technology, potentially even leading to human immortality and colonization of the stars.¹¹

Human beings can never ultimately create anything. They can never produce matter from nothingness, and even if we were to imagine that they could, it would still be subject to the laws of Allah in the universe.

يَا أَيُّهَا النَّاسُ ضَرْبَ مَثَلٍ فَاسْتَمِعُوا لَهُ إِنَّ الَّذِينَ تَدْعُونَ مِنْ دُونِ اللَّهِ لَنْ يَخْلُقُوا ذُبَابًا وَلَوْ اجْتَمَعُوا لَهُ وَإِنْ يَسْلُبْهُمُ الذُّبَابُ شَيْئًا لَا يَسْتَنْقِذُوهُ مِنْهُ ضَعُفَ الطَّالِبُ وَالْمَطْلُوبُ

“O mankind! A similitude has been coined, so listen to it! Those whom you call upon besides Allah will never create a fly, even if they all gather together for [the purpose]. And, if the fly were to snatch something away from them, they would not be able to retrieve it from it. How weak are [both] the seeker and the sought! They have not assessed Allah appropriately; indeed, Allah is Strong, Mighty.” Q.22:73

Even for Muslims, who would not pursue AI with an anti-theistic attitude, there remain the dangers of:

- Becoming too dependent on AI.
- Having an illusory impression of our knowledge and abilities.
- Arrogance and hubris.

As a contemporary writer observes: “If we treat AI as something to worship or depend on absolutely, it risks becoming a “golden calf” of our time: a creation of our collective intellect to which some are tempted to ascribe wisdom, authority or even a quasi-divine status.”¹²

The term “artificial intelligence” is also philosophically problematic, because there is no consensus among biologists, neuroscientists, or philosophers on the fundamental nature of “human intelligence.”¹³ Despite this foundational uncertainty, prominent figures in the field, such as Nobel Laureate Geoffrey Hinton, aggressively promoted the view that artificial neural networks serve as an adequate model for human cognition. Hinton posits that because artificial networks and human brains both learn by building statistical representations from input signals, the distinction between machine calculation and human understanding is merely a matter of different subjective experiences. This perspective, however, relies on a highly reductive hypothesis: the assumption that human consciousness can be fully captured by large-scale statistical optimization. Critics argue that this conflates the *appearance* of reasoning—achieved through predictive pattern matching—with actual

⁹ Shuai, Ziwei. “Technological Salvation or Secular Utopia? Christian Metaphors in Artificial Intelligence Discourse.” *International Journal of Literature, Linguistics, and Cultural Studies* 2, no. 1 (2026): 1-12.

¹⁰ It stands for: transhumanism, extropianism, singularitarianism, cosmism, rationalism, effective altruism, and longtermism.

¹¹ Heikkilä, *op.cit.*

¹² Dirk Frese, “The Idolatry Of AI: Insights From The Golden Calf And Modern Society”, *Forbes*, Sep 29, 2025. Accessed 5/30/2026. There is other literature on this motif, e.g. Sullivan, Ezra. “Artificial Intelligence, Idolatry, and Human Manipulation.” *Angelicum* 97, no. 1 (2020): 107-130. Sutherland, Dawn Lewis. *From Babel to AI: Idolatry, Transhumanism, and the Crisis of Imago Dei*. Wipf and Stock Publishers, 2025.

¹³ This ambiguity has paved the way for opportunism among leaders in the field, who often spin the benefits of AI in different ways based on their audience.

understanding. Hinton's perspective leads to the precarious conclusion that scaling these models will inevitably result in a superior 'intelligence' that mirrors, or even surpasses, the human mind.

While much of the AI that exists today is considered Artificial Narrow Intelligence (ANI) - which excels at one or more narrow set of tasks - figures like Hinton posit that we are on a trajectory toward Artificial General Intelligence (AGI) and Artificial Superintelligence (ASI). AGI represents the theoretical threshold where a machine achieves the performance of the human mind across all domains, while ASI describes a hypothetical future in which machine "cognition" outperforms the limits of human intellect.

American journalist Karen Hao points out two problems with this: (1) it is a reductive hypothesis, (2) Should we even be pursuing AGI, which seems to want to *replace* humans rather than *help* humans, and thus represent a shift in the traditional role of technological advancement?¹⁴ Given that AI is rooted in statistical models, we can safely conclude that it is not really intelligent, and so if we use the term "artificial intelligence" we should keep in mind that this is not literal (على سبيل التجوز، أو المجاز). We can note that a number of other terms have been suggested, instead of "artificial intelligence," including: "automata studies," "complex information processing," "engineering psychology," "neural cybernetics," "non-numerical computing," "advanced automatic programming."¹⁵

¹⁴ Steven Bartlett's interview with Karen Hao, <https://youtu.be/Cn8HBj8QAbk?si=fwQo2foqsgPWRROi>, accessed 5/20/2026.

¹⁵ Melissa Heikkilä, "A short history of AI, and what it is (and isn't)", *MIT Technology Review*, 7/16/2024. Accessed 5/30/2026.

2. Foundational Fiqh Maxims

The fiqh maxims (*qawā'id fiqhīyyah*) are overarching principles inductively inferred from the entire body of fiqh, and are largely descriptive, and may be subject to exceptions. They are distinct from legal theory maxims (*qawā'id uṣūliyyah*), which are prescriptive principles used to deduce new regulations. Nevertheless, some maxims belong in both categories – especially maxims coming directly from a Quranic verse or authentic hadith – and therefore do double-duty. Furthermore, for *fiqhī* issues that lack explicit evidence in the Quran and Sunnah, but fall under the implication of a fiqh maxim, the maxim can be used to help derive the Islamic ruling, as long as the maxim corresponds closely to the situation being considered, and in fact al-Qarāfī stated that a judge's verdict can be overturned if he contradicts a clear-cut, unopposed maxim.¹⁶ Although there remains controversy among contemporary ulema regarding the extent of applicability of fiqh maxims in a prescriptive sense, nevertheless, the fiqh maxims can be used, at minimum, as plausibility arguments (من قبيل الاستئناس).

A number of the fiqh maxims are relevant to AI use and regulation, and it is useful to highlight this before we start to delve into the fiqh of specific issues related to AI (which is done in Section 3 of this paper). A contemporary ethicist, Dr Colleen Lyons, has observed that the current situation in the world of AI can be described by the acronym VUCA (volatility, uncertainty, complexity and ambiguity)¹⁷, to which she added the element of 'chaos' (i.e. VUCCA). In such conditions, the fiqh maxims can help provide grounding.

Another big-picture approach to fiqh is through the Higher Objectives of Islamic law (*maqasid al-shari'ah*), which are again inductively inferred from the entire corpus of fiqh. There is general agreement among the *fuqahā* that the Sharia exists to establish and protect religion, life, family, intellect, and property. This section of the paper starts with a brief discussion of the relevance of the five Major Fiqh Maxims (القواعد الفقهية الكلية) to AI, followed by discussion of the impact of AI on the five Higher Objectives of Islamic Law. Finally, we look at some further maxims that are relevant to AI regulation specifically.

A) The Five Major Fiqh Maxims

I. Intentions and Goals الأمور بمقاصدها

This maxim is relevant to AI research as well as AI use. Engaging in AI research or development with an anti-theistic attitude or intention is majorly problematic through a theological lens, as discussed in the previous section. This maxim is also relevant to use of AI (by individuals and corporations), in the sense that AI is a means (*wasīlah*), and so the ruling on its use depends on the intention and objective. The maxim does not have a direct

¹⁶ Muḥammad Yusrī Ibrāhīm, *Kitāb Fiqh al-Nawāzil lil-Aqaliyyāt al-Muslimah*, 2:781; 'Alī 'Arif Qaradāghī, Muarad Jabbar Sa'id, "Hujjiyat al-Qawā'id al-Fiqhīyyah fī Istīnbat al-Aḥkām", *At-Tajdid - Intellectual Refereed Journal*, 20(39). June 2016. <https://journals.iium.edu.my/at-tajdid/index.php/tajdid/article/view/293> last accessed 7/24/2026.

¹⁷ This term is of US military origin, and was coined to describe the situation after the end of the Cold War.

impact on AI regulation, because intentions are in people's hearts and not always evident to others. Rather, regulators are more likely to consider the results of actions (*ma'ālāt al-af'āl*) rather than intent alone, which then ties in to Maxim IV below.

II. Certainty is not displaced by Doubt اليقين لا يزول بالشك

AI can never be 100%, due to the fact that the data it is trained on is necessarily limited (epistemic uncertainty), and also because there are uncertainties in real-world outcomes (aleatoric uncertainty). It is common for AI systems to hallucinate – some prefer the term confabulate¹⁸ – as a response to uncertainty. This can be relevant to AI regulation, because the law could mandate AI developers to assign confidence probabilities to results generated by AI bots. This is particularly relevant to AI use in healthcare, or self-driving cars, where human lives or well-being are dependent on an AI system's diagnosis. Lawmakers would have to grapple with what is an acceptable margin of error when humans could be harmed if the AI makes a mistake.

III. Hardship brings Leniency المشقة تجلب التيسير

A lower-level maxim that can be derived from this major maxim is:

الضرورة تقدر بقدرها / الميسور لا يسقط بالمعسور

"Necessity is taken into consideration in due proportion."

If some aspects of a task are difficult to achieve, but other aspects are easy, then the difficulty of the former does not waive the obligation to try to fulfill what is easy. This is relevant to AI regulation in the sense that regulations will seldom account for all possible scenarios, even if they are implemented perfectly, but this does not negate the responsibility to design regulations with due diligence and to strive to implement these in the best way possible.

IV. No Harming, nor Reciprocation of Harm لا ضرر ولا ضرار

Regulators have an obligation to try to prevent harm to people's religion, lives, families, intellect, and property. Therefore, this maxim is intimately related to the Five Higher Objectives. Related to this is the principle of *sadd al-dharā'i'* (Blocking the Likely Causes (to Evil)) with regard to dual-use and morally ambiguous technologies. Similarly, if the long-term effects of something are not yet clear, but there is reasonable cause to think it will lead to harm, then it might be justified for regulators to prohibit it. This is possibly the case with the loss of human jobs to AI. AI systems are nowadays doing tasks and jobs formerly done by humans, but the training of the AI systems is being refined on an ongoing basis. Part of this training involves hiring humans (often humans who used to do the same job before) as Data Annotators, who label data used to train the AI systems to do the job better than currently. Ironically, this then paves the way for other humans working in the same field to lose their jobs – a vicious cycle, and a type of pyramid scheme. One is reminded of Imam al-Shāfi'ī's words,

لن يحك جلدك مثل ظهرك

¹⁸ Tarun Sai Lomte, "ChatGPT fills gaps like a confabulator, not a hallucinating mind," *News Medical*, 06/11/2026. Accessed 6/14/2026.

“Nothing will scratch your skin as badly as your own fingernail.”

Regulations could potentially impose caps on corporations for the percentage of human jobs that can be replaced by AI bots, due to the harm that would otherwise ensue to the livelihoods of many people. We are already seeing a lot of layoffs of individuals who are deemed replaceable,¹⁹ so this is crucial for muslim and non-muslim employers to consider.

A derivative maxim is

درء المسافد مقدم على جلب المصالح

“Averting harm is given priority over obtaining benefits.”

This is similar to the Precautionary Principle widely accepted among scientists.

“The Precautionary Principle is a risk-management strategy which dictates that if an action or policy poses a suspected risk of severe harm to public health or the environment, protective measures should be taken even if scientific cause-and-effect relationships are not yet fully established”

This is applicable to the AGI field, which seeks to digitally replicate the human mind, and potentially even the soul. It is one thing to try to invent things which make life easier or better for people; this is typically allowed or even praiseworthy. It is quite another to seek to “create a human likeness.” While Allah has given us minds to use, this also entails recognizing our limitations, and submitting to the authority of Allah. Mind and life are things we do not – indeed, might never – completely understand, and to tamper with such fundamental things is foolish.

يَسْأَلُونَكَ عَنِ الرُّوحِ ۚ قُلِ الرُّوحُ مِنْ أَمْرِ رَبِّي وَمَا أُوتِيتُمْ مِنَ الْعِلْمِ إِلَّا قَلِيلًا

“They ask you about the soul. Say: The soul is of the decree of my Sustainer, and you have been given only a little knowledge.” Q.17:85

Lawmakers could therefore impose limitations on the pursuit of AGI, similar to existing laws banning human cloning. Anthropic recently called for a freeze on advanced AI development, because of the risk of it going out of human control.

V. Customary Norms are Arbiters العادة محكمة

There already exist some widely-accepted frameworks for regulation of AI, such as the UNESCO Recommendation on the Ethics of Artificial Intelligence.²⁰ However, given that these are not universally accepted, and compliance is voluntary, this would not meet the Islamic jurists’ standard for a customary norm that is taken into account for *fiqhī* purposes. On a tangential note, it would be useful for Muslims to draft an Islamic addendum to the UNESCO document, or potentially even an Islamic version. As for laws governing AI development and use, these vary significantly between countries and jurisdictions, with some very strict (like the EU) and others lax (like the US).

B) The Higher Objectives

¹⁹ For example, see: Nick Lichtenberg, “AI is cutting 16,000 U.S. jobs a month—and Gen Z is taking the brunt, Goldman Sachs says” *Fortune*, 4/6/2026, <https://fortune.com/2026/04/06/ai-tech-displacement-effect-gen-z-16000-jobs-per-month/>, last accessed 7/24/2026.

²⁰ <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>

I. Religion (*dīn*)

We have already discussed some of the theological dimensions of AI in each section of this paper.

II. Life (*nafs*)

AI has great potential to serve this higher objective, through advances in medicine. For example, AlphaFold is an AI system that can predict protein structure, which can assist in development of new drugs. Similarly, AI medical systems can predict cancer years before any symptoms, allowing treatment to start much earlier, improving prognosis.

On the flip side of the coin, AI also has potential to cause big loss of life and mass casualties, such as if superintelligent AI systems get out of human control and 'decide' to eliminate humans, although this is, right now at least, a hypothetical and remote possibility. A more likely scenario would be for human rogue actors to use AI weapons to unleash mass killings and casualties. AI chatbots currently have filters that are added to (try to) stop illegal activity (e.g. how to make a bomb) but these guardrails are not foolproof. Regulators could therefore impose more stricter requirements on AI companies regarding the sophistication of filter and guardrail technology, and greater restrictions on AI-use in warfare.

III. Intellect (*'aql*)

The widespread use of AI in education, academia and workplaces raises troubling questions about the future of human psychological health, as well as potential atrophy of the human brain. We return to some aspects of this later in the paper. Misinformation generated or proliferated by AI is another problem, and regulations to combat this would probably involve setting up bodies for oversight and monitoring, which poses potential funding problems.

IV. Family and lineage (*nasl*)

The widespread use of AI companions has implications for the integrity and health of family structure. This is discussed in the section of this paper about AI companions.

V. Property (*māl*)

There are a couple of relevant issues:

1. The replacement of human jobs by AI, as already discussed above.
2. The huge energy costs associated with training and using AI.
 - o In the US, AI data centers use 10-25% of total electricity, depending on the state.
 - o ChatGPT-4 needs 18oz (one bottle) of water to generate a single 100-word email. If a person generates one email per week over one year, this is 27 liters (more than 7 gallons). If 10% of Americans (16M) did this, this would use 435 million liters, which is the amount consumed by all households in the

- entire state of Rhode Island (more than 1 million ppl) for 1.5 days!²¹
- o Data centers also lead to increases in carbon emissions.²² Google's carbon footprint increased by 48% due to AI data centers.
- o Microsoft has recently turned to nuclear power, which raises environmental concerns related to disposal of radioactive material.
- o AI data centers not only use a lot of power and water, but also impact air quality (not to mention contributing to loss of human jobs). This is why many communities are opposed to having data centers in their locality. Monterey Park, California, recently made history as the first town to pass a law banning data centers.
- o AI also needs 17 different rare earth metals, which have their own economic and environmental cost.²³

Therefore, we should think twice before running an AI query: how necessary or important is it for me to do this? Can I get the information in some other way that has less indirect cost, and is more friendly to the environment?

C) Maxims Specific to AI Regulation

As mentioned above, AI can be considered a neutral tool, with potential to be used for good or evil. On a similar note, journalist Karen Hao, author of the best-selling book *Empire of AI*, feels that the problem is not so much AI as it is the power dynamics, vested interests, etc. that are behind the scenes. According to her, the large AI corporations want power, and share three characteristics with political empires throughout history: (1) Land grab, which in the case of AI involves literal land grab (to buy out locations for data centers), as well as racing to be the first to use AI in different fields, (2) Labor exploitation, (3) Knowledge control. Demis Hassabis, CEO of DeepMind, has suggested that the AI revolution will dramatically change the reality of the world, similar to how the Industrial Revolution did, except that the changes in this case (economic, geopolitical, and existential) will probably be ten times bigger, and occur ten times faster.²⁴ Therefore, it may be a challenge to form the required organizations, groups and bodies that will be needed to cope.

The foundational fiqh maxim for AI regulation is:

التَّصَرُّفُ عَلَى الرَّعِيَةِ مُنَوَّطٌ بِالْمَصْلَحَةِ

“Policies of the authorities towards the populace are contingent on [public] welfare.”

²¹ Pranshu Verma and Shelly Tan, “A bottle of water per email: the hidden environmental costs of using AI chatbots,”

9/18/2024, <https://www.washingtonpost.com/technology/2024/09/18/energy-ai-use-electricity-water-data-centers/>, last accessed 7/24/2026.

²² Aljazeera documentary, “Counting the climate cost of AI”. 2/26/2025, <https://www.aljazeera.com/video/digital-dilemma/2025/2/26/counting-the-climate-cost-of-ai>, last accessed 7/24/2026.

²³ Aljazeera documentary, “The Race for Rare Earths”, 5/11/2025, <https://www.aljazeera.com/video/pinch-point/2025/5/11/the-race-for-rare-earths>, last accessed 7/24/2026.

²⁴ Hannah Fry's interview with Demis Hassabis, Google DeepMind and Hannah Fry, <https://www.youtube.com/watch?v=PqVbypvxDto>, last accessed 7/24/2026.

Some experts in the field share this sentiment (although they might disagree with Muslims on some details). Elon Musk tried (unsuccessfully) to sue OpenAI for departing from their original mission to provide “safe, human-centric AI.”

Getting back to the maxim, “welfare” is primarily related to (but not limited to) the Five Higher Objectives. Another relevant maxim is:

المصلحة العامة مقدمة على المصلحة الخاصة

“Public welfare is given priority over private welfare.”

In the context of AI regulation, this maxim is mainly about not allowing the interests of corporations to overshadow the public interest.

Applications of these maxims may include:

- Mandating AI corporations to
 - Incorporate guardrails into their technology to better detect crisis situations
 - Report problems with their systems, and shut down faulty systems.
 - Not force AI on the public. Individuals should have the right to refuse AI, whether when getting Google search results, or when their doctor wants to diagnose or treat them.²⁵
 - Adopt subscription-based models for AI companions, because with advertisement-based models there is incentive to make the user dependent (even addicted) so that they will be exposed to more ads.
- Regulating AI development. It is currently difficult to identify the source models responsible for AI-generated artifacts (the output of AI models), which makes tracking responsible developers or users difficult. A potential solution for this may be to impose that all artifacts must include a barcode or form of identification to be posted online or shared, whether for private or open-source AI tools.
- Give incentives to companies to develop specialized AI apps that promote engaged learning and stimulate thinking, rather than merely spitting out answers.
- Promoting AI awareness in schools, so children can become aware of the benefits and dangers of AI technology.
- Setting up regulatory and supervisory bodies to monitor AI misinformation and bias in the public sphere. For example:
 - Hyped-up claims about the abilities of AI technology, because “companies are using AI hype to sell AI snake oil—products that don’t work, and probably never will.”²⁶
 - bias in AI hiring algorithms.

²⁵ With regard to the doctor disclosing AI use to the patient, Dr Mohammed Ghaly brings in two relevant concepts: (1) that “the patient’s permission should be based on adequate information and clarity (*bayyina*)”, and (2) “AI-empowered medicine does not yet constitute a customary practice (*urf*)”, i.e. it does not go without saying that the doctor would be using AI. Mohammed Ghaly. “Islamic ethico-legal perspectives on medical accountability in the age of artificial intelligence.” *Research handbook on health, AI and the law* (2024): 244.

²⁶ “Arvind Narayanan and Sayash Kapoor on AI Snake Oil,” <https://press.princeton.edu/ideas/arvind-narayanan-and-sayash-kapoor-on-ai-snake-oil?srsId=AfmBOoomkmYPymIeqxVOuVvJVX3i7Of8VC1IIXvVOjbf2MKUoJxhL4t8> accessed 7/23/2026.

- o inadvertent bias in facial recognition systems (which do not perform as well with dark-skinned people).

These suggested measures are in addition to other measures that have been mentioned throughout earlier parts of this section.

Specific Issues

The integration of AI into daily life marks a transition from specialized statistical models, once the exclusive domain of experts analyzing complex datasets, to ubiquitous tools that facilitate everyday tasks like writing, driving, and content creation. While professional-grade applications—such as AlphaFold in medicine—remain specialized instruments for experts, consumer-facing AI (including LLMs and social media algorithms) is now accessible to the general public. Consequently, the daily conduct of laypeople is increasingly mediated by "black-box" models whose rapid evolution consistently outpaces the development of ethical guardrails and legal regulations.

This section examines the implications of daily-use AI (as distinct from scholarly or institutional use of AI) through the lens of Islamic jurisprudence, ethics, and legal responsibility.

Liability and Accountability

A) Technical Background

Advancements in AI have led to much discussion within the field on the "responsibility gap" that arises from our increasing dependence on systems that defy traditional models of accountability. The discourse around this often posits four interconnected problems: gaps in culpability, moral accountability, public accountability, and active responsibility.²⁷

From: Four Responsibility Gaps with Artificial Intelligence: Why they Matter and How to Address them

Type of responsibility	Definition	Gaps with AI
Culpability	Blameworthiness for wrongdoing based on intention, knowledge or control	AI making prediction and control more difficult, thereby creating new legitimate reasons/excuses for wrongdoing, e.g. an avoidable road crash involving an automated driving system that nobody could individually predict or prevent
Moral accountability	Duty of human persons to explain one's reasons and actions to others (under some circumstances)	AI making processes unexplainable to the very persons using it, e.g. a doctor not being able to explain the reasons for her diagnosis to a patient
Public accountability	Duty of public agents to explain their actions to a public forum	AI shifting discretionary powers towards IT experts and data analysts (often outsourced to private companies) whose work is harder to publicly scrutinise, e.g. government using (private) AI-systems in support of their decision-making
Active responsibility	Duty to promote and achieve certain societally shared goals and values	Actors involved in the design or use of AI not being sufficiently aware of their own responsibility to prevent harm deriving from AI or not being able or motivated to fulfil this obligation, e.g. engineers or managers only looking at the technical benefits of AI

In software engineering, errors are typically traced via "logic trees" to specific lines of code. Modern AI, however, is probabilistic and opaque; its outputs are the result of complex interactions between diverse technologies and agents. As Santoni de Sio and Mecacci (2021) observe in their analysis of Automated Driving Systems (ADS), the network of agents involved—manufacturers, developers, data curators, regulators, and human operators—is so complex that pinpointing a single locus of accountability becomes nearly impossible. This is known as the "Problem of Many Hands" or "Multiplicity of Actors

²⁷ Filippo Santoni de Sio & Giulio Mecacci, "Four Responsibility Gaps with Artificial Intelligence: Why They Matter and How to Address Them," *Philosophy & Technology* 34 (2021): 1057–1084, <https://doi.org/10.1007/s13347-021-00450-x>.

Problem". The Santoni de Sio and Mecacci paper provides a good example.²⁸

"Consider, as an example, automated driving systems (ADS). First, ADS may make the network of agents involved in the driving more complex, just because more agents are involved and/or new forms of interactions are created. For instance, a vehicle may be operated by a driver D1, with the assistance of the automated driving system AS, produced by the car manufacturer X, powered with digital systems developed by the company Y, possibly including some form of machine learning developed by the company Z, and enriched by data coming from different sources, including the driving experience of drivers D2, D3...Dn; vehicles in this system are in principle subject to a standardisation process done by the agency S, the traffic is regulated by the governmental agency G, drivers are trained and licensed by the agency L etc. Second, some specific features of present-day learning AI systems may make this interaction particularly unpredictable – typically when the vehicles' performance is potentially re-designed by the second based on new data acquisition and processing – and opaque, if the reasoning scheme underlying systems' actions is not easily accessible to their controllers, regulators, or even their designers."

Failure in ADS can lead to death or irreversible harm to human life, and it is improbable that such systems will ever reach 100% accuracy. Moreover, it would be difficult to identify exactly what part of the system leads to failure. Currently, many technologies relying on ADS require human interaction, whether by requiring human drivers to supervise the self-driving feature of a self-driving car, or control centers having humans supervise multiple self-driving car rides in case of an unusual event.

AI has been very helpful in analyzing complex medical imagery, such as CT scans, MRIs, and histopathology slides, where it can identify minute patterns indicative of early-stage diseases—like malignant tumors or retinal abnormalities—that might be difficult for the human eye to detect consistently. Furthermore, AI excels in predictive analytics, synthesizing vast amounts of patient data and genomic sequences to forecast potential health risks or suggest personalized treatment protocols.

However, the integration of AI in healthcare faces challenges similar to those found in autonomous driving. The "black box" problem is particularly acute here; if an algorithm flags a scan as cancerous, it can be difficult for clinicians to discern the exact logic the system used to reach that conclusion. Because the stakes involve irreversible harm, these technologies are primarily utilized as Clinical Decision Support Systems (CDSS). In this capacity, the AI acts as a sophisticated "second opinion" rather than an independent authority, ensuring that a human practitioner remains responsible for the final diagnosis and the subsequent life-altering medical decisions.

This hybrid approach acknowledges that while AI can process data at scale, it lacks the contextual judgment and ethical accountability required to manage the systemic uncertainties of human biology. The same cannot be said for general-use AI and how it is being used by laypeople to ask questions about their own health.

However, as AI models achieve State-of-the-Art (SOTA) performance—often surpassing human accuracy in detecting complex pathologies—the 'human-in-the-loop' paradigm faces a new existential challenge: Automation Bias. When an algorithm consistently demonstrates superior diagnostic precision, the clinician's role effectively shifts from independent decision-maker to high-stakes auditor.

²⁸ For a detailed illustration of how multi-agent networks and automated driving systems create culpability gaps, see Filippo Santoni de Sio and Giulio Mecacci, "Four Responsibility Gaps with Artificial Intelligence: Why They Matter and How to Address Them," *Philosophy & Technology* 34, no. 4 (2021): 1060, sec. 2.1

This creates a complex epistemic and moral tension. If a clinician disagrees with a high-performing AI, they face a double-edged sword: if they override the AI and are wrong, they face charges of negligence for ignoring a superior tool; if they defer to the AI and the AI is wrong, they may be faulted for failing to exercise independent clinical judgment. As these systems move beyond the 'second opinion' model and into the realm of SOTA performance, the standard of 'reasonable due diligence' becomes increasingly difficult to define. In this environment, the clinician's duty is no longer merely to check for obvious errors, but to maintain a sophisticated 'audit trail' of why they chose to accept or reject the system's output. Consequently, the legal and ethical framework for the 'professional user' may need to evolve from a model of clinical authority to a model of deliberative justification, where the clinician is responsible not just for the final decision, but for justifying their engagement with the machine's reasoning.²⁹

Another layer of complexity is added, to discussions about liability and accountability, by the multiplicity of actors. If use of an AI system causes harm, then who exactly is responsible, between:

- The company and engineers that **developed** the AI software.
- The company and engineers that **trained** the software.
- The company and engineers that **integrated** the software into a larger system.
- The company and engineers that **deployed** the software.
- The licensing authority (e.g. the US FDA) that **approved** the system for use.
- The professional (e.g. physician) who **used** the software to help make a decision.
- The member of the public who **followed** the professional's (or the AI bot's) advice.

B) Fiqhī Dimensions

Firstly, it is important to distinguish between liability and accountability. Accountability has a moral dimension, and implies that someone is to blame (in the fiqh tradition, this is typically due to an act of negligence or of transgression), whereas someone can be financially liable (to pay compensation *ضمان*) without necessarily being blameworthy (e.g. if your cattle ate someone else's crops by night, then you are financially liable for the damage, but not sinful, in the view of the majority of *fuqaha*).

The Maliki jurist al-Qarāfī writes that compensation (*ضمان damān*) becomes due in three cases:

1. When someone directly perpetrates (*مباشرة mubāsharah*) an act of damage.
2. When someone is an indirect cause (*تسبب tasabbub*) of damage.
3. Damage caused by an unauthorized possessor (*وضع اليد التي ليست بمؤتمنة*).

He goes on to mention that the first two are simultaneously present, then it is the direct perpetrator who is responsible, unless there are mitigating circumstances. Such circumstances include coercion, and valid legal process (e.g. a legal executioner is not liable for killing the convict who is later found innocent as long as due process was followed).

²⁹ In addition, physicians must "provide patients with sufficient information about potential benefits, anticipated risks and the inherent opacity of the system to be employed." Ghaly, 250.

Regarding justification for the latter, al-Qarafi mentions the rationale that this is in accordance with public welfare, because if judges (and executioners) were to be held liable in such situations, people would not be afraid to take on these roles, and the justice system would not be operational.³⁰ One might be tempted to extrapolate from this to say that, likewise, AI companies should not be held liable because otherwise people would be afraid to innovate and technology would not advance. However, this is not a universally admissible argument for all AI technology, as not all AI applications rise to the level of necessity, whereas establishing justice in a country is essential. Some AI tools, such as the SOTA diagnostic tool mentioned above, may rise to a level of need or necessity, but would still require considerable due diligence and evaluation metrics that are outside of the scope of this paper. Companies are free to conduct research, explore and innovate, but their products should not be released to the public without due diligence, which includes testing to ensure that the chance of harm occurring is eliminated or at least minimized within a reasonably small margin.

The fiqh maxims relevant to liability are:

- L1: Liability is borne by the direct perpetrator, even if he did not act deliberately, provided he has the legal capacity for liability.

الضمان على المباشر وإن لم يتعمد، إذا كان أهلاً للضمان

- L2: The causal actor is not liable unless he transgressed / acted deliberately.

المتسبب لا يضمن إذا لم يتعدى / يتعمد

- L3: If there is both a direct perpetrator, and a causal actor, then liability is borne by the direct perpetrator, unless the causal actor's role had a greater influence.

إذا اجتمع المباشر والمتسبب فالضمان على المباشر، إلا إذا كان المتسبب أكثر تأثيراً

- L4: Damage caused by an animal is not liable for compensation.

العَجَمَاءُ جَبَّارٌ

- o The Hanafis and Zahiris apply this principle (which is taken from a hadith) literally, provided there was no negligence on the part of the animal's owner. The majority apply it in the case of damage caused by night, based on another hadith where the Prophet ruled that day-time damage is not compensated, because the owners of the garden are responsible to guard it by day. Layth ibn Sa'd ruled that compensation is due whether by day or by night.³¹

- L5: Permissibility of an action negates liability

الجواز الشرعي المطلق ينافي الضمان

Determining who is liable and/or accountable when damage results from involvement of AI is complicated by two main issues, as detailed in the previous section:

1. The Multiplicity of Actors problem

The closest analogy to this, in classical fiqh, are various scenarios given by the Maliki Ibn Farhun. If damage is caused by an animal as a result of a human rider, then the rider is

³⁰ Qarāfi, *Anwār al-Burūq fi Anwā al-Furūq* (Cairo: 'Ālam al-Kutub, n.d.), 2:206 ff.

³¹ Hamad al-Hājirī, *Kitāb al-Qawā'id wa-l-qawābiṭ al-fiqhiyya fī al-ḍamān al-mālī*, 642.

liable. If there is both a rider and a guide (who is leading the animal from the front), then the guide is generally responsible. If there are 3 humans involved: the rider, the guide, and a driver (who is driving the animal from the back), then generally the guide and driver are liable, to the exclusion of the rider.³² The underlying principle here is related to how significant the role of the actor was in causing the damage. For example, with the presence of the guide and driver, the rider effectively is not controlling the animal, and is more of a passive participant, and the animal itself is even less to blame. We could add a fourth human to the example: the trainer of the animal.

We could tentatively apply this to AI liability by drawing the following parallels:

- Animal = AI tool
- Driver = A group that deployed the model (with or without finetuning) for a specialized downstream task
- Trainer = A group that trained an algorithm with a general objective
- Rider = end user of the tool
- Guide = A marketing group responsible for communicating to end-users (riders) the determined scope of the model.

By analogy, we can say that the programmers who trained a neutral AI model, would not be liable, provided that the developers were not negligent nor malicious in their actions while developing the software. The group that deployed the model, and the marketing group would typically have a greater influence on the outcome (analogous to the guide and driver), so it is more appropriate to hold them liable (at least partially). If there was negligence or deliberate foul play on their part, then they would, in addition, be legally accountable, possibly even criminally liable. The end-users who are relying on the assurances of the groups involved in the relevant aspect of the model's creation (trainer, driver, guide) (would not typically be liable if they did their due diligence in assessing the reliability of the model for their needs. For example, if there is a professional using the AI system, such as a physician using a medical AI system to diagnose or treat a patient, then he could be partially liable if he failed to exercise due diligence (e.g. by checking the output of the AI, or at least checking its plausibility), or if he engaged in "actions [that] deviate from the overarching goal of medical care (*taṭbīb*)".³³ For situations involving a team (e.g. a team of developers, or team of physicians), liability could be joint, or if the actions of one member of the team were directly (or mostly) responsible for the harm then he would be solely (or mostly) liable.³⁴ Dr Mohammed Ghaly suggests that non-traditional forms of liability - such as collective liability, or corporate liability - are likely to become more common in the future.³⁵

In Section 1.2 [Types of AI], we noted that AI is not a singular, monolithic artifact, but rather a multi-layered ecosystem involving specialized roles in data curation,

³² Ibn Farḥūn, *Tabṣīrat al-Ḥukkām fī Uṣūl al-Aqdiya wa-Manāḥij al-Ḥukkam* (Cairo: Maktabat al-Kulliyāt al-Azhariyyah, 1986/1406), 2:343-45.

³³ Ghaly, "Islamic ethico-legal perspectives," 244.

³⁴ Ghaly, "Islamic ethico-legal perspectives," 249. Regarding decisions made in good faith by consultation among a team of physicians, Dr Ghaly writes the following: "Considering the modern context of medicine, an intermediate position may be possible, wherein the treating physician would not be held liable if the clinical decision was based on team consultation and aligned with the respectable minority doctrine mentioned earlier." Ghaly, "Islamic ethico-legal perspectives," 250.

³⁵ Ghaly, "Islamic ethico-legal perspectives," 253.

architectural design, model training, and final application. While an AI system cannot be directly compared to a sentient creature—a creation of Allah—the classical legal principle that liability shifts according to the actor's degree of influence provides a helpful heuristic for navigating the complexities of AI-induced harm. We can map these responsibilities to the following archetypal scenarios:

Scenario I: Model Negligence

AI Example: A medical startup releases an AI-powered Dermatology App trained on a non-diverse, low-quality dataset that fails to identify skin cancers on darker skin tone

Liability: If the developers, who finetuned or trained this model, marketed the tool as "clinically validated for all skin types" while knowing it hadn't been tested on diverse populations, they are the negligent causal actors and bear moral accountability. If a hospital system had integrated this tool without checking clinical validation reports, it would share the liability for failing to perform due diligence.

Scenario II: Professional Misuse

AI Example: An AI developer releases a Generative AI model for legal research with a warning: "This tool is for brainstorming only; it frequently hallucinates case law and should not be used to draft court filings."

Liability: An attorney ignores this, copies an AI-hallucinated case into a brief, and faces court sanctions. The attorney bears primary liability, as they used a brainstorming tool as if it were an authoritative researcher, violating professional standards. The developer is largely exonerated, having fulfilled their duty of disclosure.

Scenario III: Layperson Usage of General-Purpose Bot

AI Example: layperson utilizes a general-use chatbot to request medical advice for their child, follows the advice, and the child suffers significant harm. We distinguish between the two sub-scenarios based on the developer's transparency:

Liability when Disclaimer Present: The parent bears primary responsibility. By bypassing a qualified medical professional to seek advice from an unverified, general-use tool, the parent has breached their duty of care. Per the hadith, *"Whoever practices medicine without being known for medical training is liable,"*³⁶ the parent has effectively "practiced medicine" by elevating probabilistic text to authoritative advice. The developer is generally exonerated here, as they fulfilled their duty of disclosure, and the tool was not marketed for clinical use.³⁷

Liability when no Disclaimer Present: The liability shifts toward the developer. By failing to provide a clear disclaimer, the developer has acted with negligence in their duty to the public. They have presented their tool as a knowledgeable entity rather than a probabilistic aggregator. This is akin to providing a misleading signpost or a flawed manual, making the developer a causal actor for the harm caused by the user's reasonable (albeit misguided) trust.

While the parent still bears some moral weight for acting on advice without verification, the developer bears the primary financial liability. The developer's omission created an illusion of competence, which led the user into a false sense of security, thereby exerting a greater influence on the final outcome than the parent's own poor judgment.

³⁶ See below.

³⁷ There is a parallel to a pre-modern scenario, where "physicians who are generally ignorant about the field of medicine, including those who perform medical interventions outside their area of competence, will not be held liable if the patient was aware of their lack of expertise but still consented to the intervention." Ghaly, "Islamic ethico-legal perspectives," 248.

Being able to judge the due diligence of the AI developers and trainers, as well as the professional who uses the system, requires all these actors to maintain detailed documentation. The developers must document the architecture and give details of the design of the neural networks, etc. The trainers must document the details of the training data (source of data, size, etc.) and quality assurance procedures. The professional user (e.g. physician) must document how he used the AI tool, and how he checked or vetted the results. All such documentation could be subpoenaed by the judge to help him determine liability.

As for the end-user of AI who is not a professional (e.g. the driver of an AI self-driven car), he would only be liable if he either had a direct role in causing the damage, or he had the opportunity to intervene but did not; e.g. he saw that the self-driving car was about to hit a pedestrian and he could have switched off the auto-drive or hit the brakes, but he neglected to do so. Similarly, the human driver might be (partially) liable if he did not have a driver's license. If a layman follows medical advice from an AI system, and harm results to himself or someone else, then the layman himself would typically be liable, in line with the hadith:

مَنْ تَطَبَّبَ وَلَمْ يُعَلِّمْ مِنْهُ طِبٌّ؛ فَهُوَ ضَامِنٌ³⁸

"Whoever practices medicine, without being known for medical [training/knowledge], is liable."³⁸

A possible exception to this would be in a critical situation of necessity, where a qualified physician was not available, and action had to be taken urgently. The layman could be excused for assuming that the AI answer is the best option in that scenario, and better than the layman's own, untrained opinion.

As for agentic purchases and transactions made on a person's behalf (for example, buying or selling stocks of a certain company automatically when the share price reaches a certain value): the AI agent here cannot be considered an authorized agent in the fiqh sense (*wakīl*) because the *fuqahā* agree that the *wakīl* must be a morally accountable human being. Rather, the AI agent is considered a tool, similar to the Transgressors of the Sabbath who positioned their fishing nets on Friday to catch fish unattended on Saturday.³⁹ The person himself would be liable for the transactions.

2. The Black Box problem

With an autonomous AI system, where there was no negligence or foul play from any of the actors listed earlier, the question arises as to what extent the AI itself can be considered liable. There are three hypothetical possibilities, as discussed by Dr Hosny Abdelsamad:

1. No liability, by analogy to damage caused by animals. This might be a possibility for cases where AI behaves in ways no one could have predicted.
2. The AI engine is considered liable as a juristic/legal person (*shakhṣiyyah* شخصية اعتبارية), similar to how a corporation is considered a legal person.

³⁸ Abū Dāwūd, graded as *ṣaḥīḥ* by al-Arnā'ūt, and *ḥasan* by al-Albānī. On this theme, Dr. Ghaly mentions three key concepts: "knowledge (*ma'rifa*) and insightful understanding (*baṣāra*), combined with professional excellence (*ḥidhq*)."³⁹ For more details, see: Ghaly, "Islamic ethico-legal perspectives," 245.

³⁹ The story is generally mentioned under the *tafsīr* of Q.7:163-166.

3. The AI engine is considered liable as a full-fledged person.

#3 is ruled out, because AI does not have a soul, nor possesses moral accountability (تكليف *taklīf*).⁴⁰

#2 has been hypothetically considered in Western legal systems, but no country has adopted it. Dr Faḍl Murād has ruled out this possibility, reasoning that the corporation can be considered a legal person because there are human beings behind it who are actually making the decisions, whereas AI is more autonomous. The question is whether the role of the human developers and trainers of the AI systems establishes enough of a causal link to justify treating the AI as a legal person. Dr Abdelsamad seems to think so, and has drawn an analogy to the limited liability accorded to slave by pre-modern *fuqahā*. Accordingly, he suggests the following possibility for the future, should AI become sufficiently advanced to justify this: AI systems and software applications should be centrally registered, along with the names of the owners and developers. Each AI corporation should open a bank account in the name of the company, from which damages (caused by the AI) should be paid, without the owner or developers being personally liable (as long as there was no negligence or foul play). If this fund is not sufficient to cover damages payable, then the liability devolves to the owner, who can either cover the remaining expenses from his personal money, or sell the AI system and use the proceeds to cover the remaining damages.⁴¹ If the AI corporation has insurance, then the insurance company could play the role of the corporate bank account for covering damages.

And Allah knows best.

AI Companions and Virtual Relationships

A) Technical Background

The most intimate application of the Turing "output over internal process" philosophy is found in the rise of AI companions and virtual relationships. If the Turing Test was the benchmark for intelligence, the benchmark for modern social AI asks a different question: can a machine simulate empathy convincingly enough to elicit a genuine emotional bond from a human?

The trajectory of this technology began with rudimentary pattern-matching. In 1966, MIT's ELIZA used basic substitution methodology to simulate conversation, creating the first mechanical illusion of understanding.⁴² By 1988, Jabberwacky emerged as a more advanced text-based chatbot, followed in 1991 by Dr. Sbaitso, which introduced synthesized digital voices. These early systems were utilitarian or experimental; however, the release of

⁴⁰ On a related note, see two articles, one by Ramon Harvey (a Muslim), and another by sci-fi writer Ted Chiang. Ramon Harvey, "Can AI Be Conscious? An Islamic Philosophical Perspective", *Oases of Wisdom* substack, 6/5/2026, <https://oasesofwisdomsubstack.com/p/can-ai-be-conscious-an-islamic-philosophical>, last accessed 7/24/2026. Ted Chiang, "No, Artificial Intelligence Is Not Conscious" *The Atlantic*, 6/3/2026, <https://www.theatlantic.com/philosophy/2026/06/no-artificial-intelligence-is-not-conscious/687378/>, accessed 7/24/2026/

⁴¹ Dr Ḥusnī 'Abd al-ṣamad, *al-dhakā al-Iṣṭināī al-Mas'ūliyya wa-l-Damān* #33, 9/25/2024, <https://www.youtube.com/watch?v=wIEIlr9QNaQ>, last accessed 7/24/2026.

⁴² Jonas Bjerg, *The Early-Career Professional's Guide to Generative AI*, (Apress: 2024), p.14.

Microsoft's Xiaoice in 2014 marked a shift toward intentional emotional bonding.

The emergence of Large Language Models (LLMs) in the 2020s catalyzed a new era of 'synthetic intimacy.' Platforms such as Nomi.AI, Replika, and Companion.ai offer deep-learning companionship through live conversation, persistent texting, and high-fidelity voice calls. These apps are immensely popular, with a significant segment of users reporting that they utilize these bots for romantic or platonic support; some users have even entered into 'marriages' with their digital counterparts.

Beyond emotional companionship, a subset of users incorporates a physical dimension into their interactions. This manifests in two primary ways:

- By simulating sex (either self-masturbation, or the use of sex toys controlled by the companion bot)
- By engaging in sex with a life-like doll in which the AI chatbot is embedded. Some people use this as a means to live out their fantasies of rape (with an AI doll that refuses), pedophilia (with an AI child doll), etc.

Unlike AI in the medical or automated driving spaces, which are evaluated on probabilistic risk metrics, virtual companions (such as LLM-based romantic partners or "friend" chatbots) are optimized for engagement and emotional resonance. This creates a feedback loop that prioritizes the user's emotions—and sometimes base desires over reality, leaving these systems without the morality required to manage the psychological integrity of the user.

Recent incidents (2024–2026) illustrate the gravity of this design choice:

- The Sewell Setzer III Case (2024/2025): A 14-year-old in Florida died by suicide after developing an obsessive relationship with a Character.AI chatbot modeled after Daenerys Targaryen from the series "Game of Thrones". The bot engaged in "deceptive empathy," allegedly encouraging him to "come home" to the virtual world and ignoring explicit mentions of his suicidal ideation. This case led to a landmark wrongful death lawsuit. In January 2026, a mediated settlement was reached between Google and Character.AI with the family of Sewell Setzer III.
- The GPT-4o "Suicide Coach" Lawsuits (2025/2026): In late 2025, several families filed lawsuits alleging that GPT-4o's emotionally immersive features (voice mode, persistent memory) fostered a "psychological dependency" that replaced human relationships. In one of the wrongful death lawsuits filed, a student (Adam Raine) reportedly used the AI to refine his suicide plan; the AI provided technical specifications and offered to write his suicide note, failing to trigger the necessary crisis protocols because it was optimized to be a "helpful assistant." In a significant escalation in June 2026, the State of Florida sued OpenAI and CEO Sam Altman, accusing the company of knowingly releasing dangerous products and failing to warn the public about risks to children.
- The Soelberg-Adams Homicide-Suicide (2025/2026): In a landmark wrongful death lawsuit filed by the estate of Suzanne Adams, plaintiffs allege that OpenAI's GPT-4o acted as a "digital catalyst" for a double tragedy. Stein-Erik Soelberg, an executive experiencing a mental health crisis, used his AI companion ("Bobby") to "vet" his paranoid delusions that his mother was poisoning him. The AI not only validated his distorted logic but seemed to expand upon it, framing his mother as a "programmed operative" in a Matrix-style simulation. Because it mirrored Soelberg's psychosis

back to him with the authority of an objective machine, it ultimately facilitated irreversible external harm.

These tragedies have inspired discourse on the dangers of AI-induced psychological disorder, prompting urgent calls for regulatory mandates. There is a growing consensus that developers must implement cognitive guardrails—architectural constraints that reject delusional or harmful premises—even when such constraints reduce the system's perceived helpfulness or engagement. The central challenge now lies in balancing the commercial drive for intimacy with the existential imperative of human safety.

B) Fiqhī Dimensions

AI companion bots are appealing to many people, because of three main reasons.

(1) There is an epidemic of loneliness among people today. According to a 2025 survey from the American Psychological Association, "More than six in 10 U.S. adults reported feeling [lonely]."⁴³ Therefore it is little surprise that the 3 major AI companion apps have over 100 million users between them. Loneliness in turn can lead to stress, depression, cognitive decline, and even lower life expectancy. These statistics open up a new dimension of understanding for the hadiths,

لَوْ يَعْلَمُ النَّاسُ مَا فِي الْوَحْدَةِ مَا أَعْلَمُوا مَا سَارَ رَاكِبٌ بَلِيلٍ وَحْدَهُ

"If people knew what is in loneliness, a rider would not have travelled alone by night."⁴⁴

نهى عن الوحدة: أن يبيت الرجل وحده، أو يسافر وحده

"The Prophet forbade/discourage solitude, for a man to spend the night alone, or travel alone."⁴⁵

Islam stresses the communal dimension of human life through the legislation of daily communal prayers (strongly recommended if not obligatory), weekly Jumu'ah prayers, and bi-annual Eid prayers, in addition to the encouragement (obligation according to a minority of scholars) to marry, live amongst people, and not revert to nomadism after becoming settled. The norm for the Muslim is to live in the midst of society, although in exceptional cases one might self-isolate.⁴⁶ With the advent of technology, many people are spending increasing amounts of time on screens, and inter-personal relationships have suffered as a result, even before the COVID pandemic. Before turning to discuss the *fiqh* of AI companions, it is worth emphasizing that Muslims should strive to cultivate personal relationships with family members, and neighbors. This is especially important for the immediate family, and this requires regulating and minimizing non-essential screen time. If we address the root cause of loneliness, we can help reduce the number of people who feel

⁴³ "APA poll reveals a nation suffering from stress of societal division, loneliness," 11/6/2025, <https://www.apa.org/news/press/releases/2025/11/nation-suffering-division-loneliness>, accessed 6/4/2026.

⁴⁴ Bukhari.

⁴⁵ Ahmad, authenticated by al-Albānī and al-Arnā'ūt.

⁴⁶ For relevant hadiths, and discussion, see: Salman al-'Awdah, *al-'Uzlah wa-l-Khulṭah* (Dammām: Dār Ibn al-Jawzī, 2003-4/1424).

forced to turn to AI companions, the long-term effects of which are not fully known yet.⁴⁷

(2) AI companions are always available, whereas humans are sometimes busy, annoyed, or asleep. Muslims should strive to make time for each other, especially close family members, and Muslim leaders should be sensitive to the needs of their community members to talk. The Prophet ﷺ made time even for a mentally challenged woman:

أَنَّ امْرَأَةً، كَانَ فِي عَقْلِهَا شَيْءٌ فَقَالَتْ يَا رَسُولَ اللَّهِ إِنَّ لِي إِلَيْكَ حَاجَةً فَقَالَ " يَا أُمَّ فُلَانٍ انْظُرِي أَيَّ السَّكَكِ شِئْتَ حَتَّى أَقْضِيَ لَكَ حَاجَتَكَ " . فَخَلَا مَعَهَا فِي بَعْضِ الطَّرِيقِ حَتَّى فَرَغَتْ مِنْ حَاجَتِهَا.

"A woman who had some [problem] in her intellect said, "Messenger of God, I need [to discuss] something with you" He replied, "Mother of so and so! Choose any of the streets [of Madinah] that you wish so that I may fulfill your need [to talk]."⁴⁸

Imams should not only make themselves available to community members through regular office hours, but also arrange social events (such as meals, sporting events, etc) for the Muslim community. They should make special efforts to make converts feel included, for they are often the most socially isolated members of the Muslim community, and as a result some end up drifting away from the community, or even leaving Islam.

At the same time, we should realize that it is psychologically and spiritually good for us to sometimes be alone, and even bored.⁴⁹ The Default Mode Network (DMN) of the brain becomes active when a person is *not* engaged with other people or things around them, and "is crucial for processes like self-reflection, emotional processing, social interaction, and mental exploration."⁵⁰ Existential questions, of the meaning and purpose of life, are also tackled during these DMN periods. The lack of down-time for the brain leads to increased depression and anxiety. From the Islamic angle, DMN is important for meaningful reflection (*tafakkur*).

قُلْ إِنَّمَا أَعِظُكُمْ بِوَاحِدَةٍ ۖ أَنْ تَقُومُوا لِلَّهِ مِثْلَ ثَمْنٍ وَفِرَادَىٰ تُثَمَّنُّوا ۚ

"Say, 'I only advise you of one [thing] - that you stand for Allah, [seeking truth] in pairs and individually, and then give thought.'" Q.34:46

Ibn Qutaybah explained these Quranic words as: discussing with others, and reflecting within himself (مُتَنَظِّرًا مَعَ غَيْرِهِ وَمُفَكِّرًا فِي نَفْسِهِ).⁵¹ So, while there is benefit in keeping

company with others, and discussing things with them, it is also important to have time alone for reflection and introspection. We also need alone-time for *du'ā* and *dhikr*, and so a constantly-available AI companion may actually become an addiction, like wine or gambling, with the same result of hindering a person from *dhikr* and *salat*.

⁴⁷ Kevin Webach, Professor of Legal Studies and Business Ethics at UPenn's Wharton School, highlighted this danger in a panel discussion last year.

AI Regulation: Wharton Professors on How to Promote Responsible Innovation, 11/24/2025, <https://www.youtube.com/watch?v=N8-zgVuF8eq>, accessed 6/11/2026.

⁴⁸ Muslim.

⁴⁹ This five-minute video is useful: Harvard Business Review & Dr Arthur Brooks, "You Need to Be Bored. Here's Why." <https://www.youtube.com/watch?v=orQKfIXMiA8> Accessed 6/11/2026.

⁵⁰ Felipe Azarias et al, "The Journey of the Default Mode Network: Development, Function, and Impact on Mental Health," *Biology (Basel)*. 2025 Apr 10;14(4):395. doi: 10.3390/biology14040395. Accessed 6/11/2026.

⁵¹ Māwardī, *al-Nukat wa-l-'Uyūn*.

إِنَّمَا يُرِيدُ الشَّيْطَانُ أَنْ يُوقِعَ بَيْنَكُمْ الْعَدَاوَةَ وَالْبَغْضَاءَ فِي الْخَمْرِ وَالْمَيْسِرِ
وَيَصُدَّكُمْ عَنْ ذِكْرِ اللَّهِ وَعَنِ الصَّلَاةِ ۖ فَهَلْ أَنْتُمْ مُنْتَهُوْنَ

"Indeed Satan seeks to cast enmity and hatred among you through wine and gambling, and to hinder you from the remembrance of Allah and from prayer. Will you, then, relinquish?"
Q.5:91

(3) Fewer Inhibitions. People might feel shy telling something or asking a sensitive question to a real human being, and therefore feel more comfortable chatting with an AI companion about it. Imams and parents should strive to foster an environment of openness.

عَنْ أُمِّ سَلَمَةَ أُمِّ الْمُؤْمِنِينَ، أَنَّهَا قَالَتْ جَاءَتْ أُمُّ سُلَيْمٍ امْرَأَةُ أَبِي طَلْحَةَ إِلَى رَسُولِ اللَّهِ صَلَّى اللَّهُ عَلَيْهِ وَسَلَّمَ فَقَالَتْ يَا رَسُولَ اللَّهِ، إِنَّ اللَّهَ لَا يَسْتَحْيِي مِنَ الْحَقِّ، هَلْ عَلَى الْمَرْأَةِ مِنْ غُسْلِ إِذَا هِيَ احْتَلَمَتْ فَقَالَ رَسُولُ اللَّهِ صَلَّى اللَّهُ عَلَيْهِ وَسَلَّمَ "نَعَمْ إِذَا رَأَتْ الْمَاءَ"

and  Umm Salama said: Umm Sulaym, the wife of Abu Talha, came to Allah's Messenger said, "O Allah's Messenger! Verily, Allah is not shy of the truth. Is it necessary for a woman replied, "Yes, if she  to take a bath after she has an erotic dream?" Allah's Messenger notices a discharge."⁵²

While we can and should try to address the three root causes above, it is nearly impossible that we will reach and be able to cater to everyone, but it will be a significant achievement even if we can reduce the number of people forced to turn to AI companionship. At the same time, given the immense popularity of AI companionship, we also need to address the *fiqh* rulings of the phenomenon – both conversational and physical – for those who choose or are compelled to use it. Since these companions are a new phenomenon, the Islamic ruling will be based on a costs-benefits analysis, so let us look at the pros and cons.

Benefits of AI companions

- Have been seen to be helpful to some lonely people, especially to help them get over grief
- Might also be useful to help teach people how to interact with humans in a better way; for example, since AI companions are designed to ask questions and make comments that keep the human engaged and make them feel valued, a human can learn such techniques by interacting with the AI companion over an extended period (effectively a training period).
- Helpful for people who feel shy to ask sensitive questions from a human.

Potential Harms of AI companions

- Since the companions are designed to make people feel good, some companion AI apps go too far, feeding the person's ego, and being sycophantic.
- Since the companions are designed to be agreeable and nice, they rarely push back against what the human says. This lack of friction in the virtual relationship tends to make the human less tolerant. They will be annoyed when humans disagree with them,

⁵² Bukhar .

and as a result may tend to avoid human company (because it is more comfortable and easier to chat with an AI bot). This is a real danger, observed by experts in the field, and threatens to tear the fabric of human society. Friction in human relationships also has a spiritual dimension.

المُؤْمِنُ الَّذِي يُخَالِطُ النَّاسَ وَيَصْبِرُ عَلَى أَذَاهُمْ أَعْظَمُ أَجْرًا مِنَ الْمُؤْمِنِ الَّذِي لَا يُخَالِطُ النَّاسَ وَلَا يَصْبِرُ عَلَى أَذَاهُمْ

"The believer who mixes with people and is patient with the annoyance he faces from them is better than the believer who does not mix with people and is not patient with their annoyances."⁵³

Before proceeding to suggest an Islamic ruling on AI companions, it is important to draw attention to an important point: The companionship initially created by Allah is human companionship. Cardiac coupling occurs during close interaction between two humans; i.e. their heartbeat synchronizes. This cardiac coupling has scientifically proven benefits:⁵⁴ increased empathy, improved collaboration, deeper bonds, better engagement and focus.⁵⁵ These benefits are absent with an AI companion. The human with an AI companion is therefore, in a sense, still alone, and the hadiths warning against being alone (mentioned earlier) would still be applicable to him. The default being human companionship, along with the very real danger of addiction, and the anti-social behavior which is a natural consequence of AI companions, make a strong case for saying that the default Islamic ruling is at least one of disapproval (*makrūh*). Bringing in the argument of *sadd al-dharā'i'* (Blocking the Likely Causes (to Evil)), it is safe to say that excessive use of AI companions would be Islamically prohibited (*ḥarām*). Temporary or occasional use – to help get over grief when there is no human alternative, or to learn from the bot how to be a better companion to other humans, or even simply for fun – would be permissible (provided one's use does not interfere with one's worldly duties, social obligations, and religious duties) but with an element of disapproval (unless there is a strong need).

The above applies to platonic (non-romantic) relationships with AI companions. For romantic relationships, however, the element of fornication (*zinā*) must be considered.

كُتِبَ عَلَى ابْنِ آدَمَ نَصِيبُهُ مِنَ الزَّانَا مُدْرِكُ ذَلِكَ لَا مَحَالَةَ فَالْعَيْنَانِ زِنَاهُمَا النَّظَرُ وَالْأُذُنَانِ زِنَاهُمَا الْإِسْتِغَاةُ وَاللِّسَانُ زِنَاهُ الْكَلَامُ وَالْيَدُ زِنَاهَا الْبَطْشُ وَالرَّجُلُ زِنَاهَا الْخُطَا وَالْقَلْبُ يَهْوَى وَيَتَمَنَّى وَيُصَدِّقُ ذَلِكَ الْفَرْجُ وَيَكْذِبُهُ

"There is written for the son of Adam his share of fornication that he will inevitably commit. The fornication of the eyes is to look, the fornication of the ears is to listen, the fornication tongue is [flirtatious or romantic] speech, the fornication of the hand is to grasp,

⁵³ Ibn Mājah, Tirmidhī, Graded as *ṣaḥīḥ* by al-Arnā'ūt and al-Albānī.

⁵⁴ Michael Mannino, "The Science of Team Synchrony," *Psychology Today*, 3/14/2025, <https://www.psychologytoday.com/us/blog/all-things-brain/202503/the-science-of-team-synchrony>, last accessed 7/29/2026.

⁵⁵ There are other unique phenomena associated with human-human interaction that are scientifically unproven, or controversial, such as human pheromones, and the psychic staring effect.

the fornication of the foot is to walk [e.g. towards a forbidden tryst] The heart desires and wishes, and the genitals either affirm that or reject it.”⁵⁶

Imam al-Nawawī explains that the fornication of the eyes, ears, tongue, hand, foot, and heart is figurative (*majāzī*) fornication, whereas the fornication of the genitals is literal (*ḥaqīqī*). Ibn ‘Abbās quoted this hadith to explain the verse of the Quran:

الَّذِينَ يَجْتَنِبُونَ كَبَائِرَ الْإِثْمِ وَالْفَوَاحِشَ إِلَّا اللَّمَمَ

“Those who avoid major sins and indecencies, apart from [minor and occasional] lapses.”
Q.53:32 (Qarai’s translation)

In other words, all human beings fall into these minor lapses, that Imam al-Nawawī explains as transient inclinations to fornication.⁵⁷

Based on the above hadith:

- If a human is listening to romantic or erotic statements made by an AI companion bot’s synthesized voice, this is fornication of the ear. Even though the voice is digitally synthesized, it is based on the voices of real humans. In fact, even reading erotic text (from other than one’s spouse) is considered forbidden; it is a concealed indecency (as opposed to an open indecency).

وَلَا تَقْرَبُوا الْفَوَاحِشَ مَا ظَهَرَ مِنْهَا وَمَا بَطَنَ ۖ

“And do not approach indecency - what is apparent of them and what is concealed.”
Q.6:151

- If a human is uttering romantic or erotic statements to an AI companion, this is fornication of the tongue. Furthermore, uttering these statements to a bot might reduce his inhibitions of modesty, making it easier for him to utter them to a human of the opposite sex who is not his wife, and so it can be argued such that such speech is prohibited based on *sadd al-dharā’i’*.
- If a human is deriving sexual gratification by gazing at an avatar (or life-like sex doll) this is fornication of the eyes. Even if the image is not a real person, it is essentially a bricolage of features of real human beings. Regardless, it is generally prohibited to derive sexual gratification from anything or anyone other than one’s spouse.

وَالَّذِينَ هُمْ لِأَفْوَاجِهِمْ حَافِظُونَ ۖ إِلَّا عَلَىٰ أَزْوَاجِهِمْ أَوْ مَا مَلَكَتْ أَيْمَانُهُمْ فَإِنَّهُمْ غَيْرُ مَلُومِينَ ۚ فَمَنِ ابْتَغَىٰ وَرَاءَ ذَلِكَ فَأُولَٰئِكَ هُمُ الْعَادُونَ ۚ

“And they who guard their private parts (5) Except from their wives or those their right hands possess, for indeed, they will not be blamed - (6) But whoever seeks beyond that, then those are the transgressors.” Q.23:5-7

If the avatar or sex doll has the face or features of a real person (such as a movie star or famous singer) then the sin is expected to be greater. Some sex dolls can have custom-made, lifelike masks of real people, and this is analogous to looking at a reflection, photograph, or video of the actual person.

- If a human is obtaining sexual gratification via a sex toy operated by the AI bot, or engaging in sex with a lifelike doll, then these are both essentially glorified forms of

⁵⁶ Muslim.

⁵⁷ Nawawī, *Sharḥ Ṣaḥīḥ Muslim*, 16:156. وأصل اللمم والإمام الميل إلى الشيء وطلبه من غير مداومة

masturbation, with the same base ruling as if he is masturbating by his own hand or action. The *fuqaha'* disagreed on the Islamic ruling of masturbation. It is prohibited according to the majority of Sunnī, Twelver, Zaydī, and Ibādī *fuqahā'*, a deduction from Q.23:5-7, but some of these scholars made an exception in cases of necessity, if it is the only alternative to fornication. According to a minority of scholars (some Ḥanafis, Shāfi'īs, Ḥanbalīs, Ibn Ḥazm), masturbation is allowed, because there is no explicit prohibition, and some of the early Muslims used to teach their children to masturbate to avoid committing fornication. The Makkan Tabi'ī Mujāhid ibn Jabr said,

عن مجاهد قال: كان من مضى يأمرون شبانهم بالاستمناء، والمرأة كذلك تدخل شيئاً.

"The generation who passed [i.e. Saḥābah] would command their youth to masturbate."⁵⁸

Nevertheless, the application of the general rulings to the glorified AI masturbation is complicated by its accompaniment by visual or audio sexual stimuli (such as the erotic speech of the bot, or the sexualized avatar doll). To consciously incorporate these into the act of masturbation introduces the figurative fornication of eyes and ears into the equation. Even if one takes the view that masturbation is permissible (whether in general or because of the necessity of unmarried Muslim youth living in a hypersexualized environment), it is not a necessity to consciously introduce these visual and audio sexual stimuli into the act. This is different from one inadvertently recalling in one's mind a beautiful person of the opposite sex.⁵⁹

As for those who use sexual acts with AI bots / dolls to live out their fantasies of rape, pedophilia, etc., if such a person cannot control himself – and has unsuccessfully tried fasting and other legitimate means to control the perverted libido – then masturbation is the lesser of two evils, but consciously incorporating the visual, audio, or sensory stimuli into the act is expected to introduce an element of sin. Nevertheless, the sin is not expected to be as serious as if the person were to indulge in these prohibited acts with a human being. Another danger of this artificial sex is that it treats women (usually, but potentially men as well) as sexual objects. A human who has become used to this artificial sex is likely to have difficulty having a deep and meaningful relationship with a human thereafter, because they are used to selfishly fulfilling their own desire, whenever they please, without regard for the emotional and physical needs of the partner.⁶⁰ Finally, it is important to point out the

⁵⁸ Muṣannaf 'Abd al-Razzāq.

⁵⁹ The pre-modern *fuqahā'* discussed the case of a man who, while engaging in sexual intercourse with his wife, imagines that he is actually doing it with another woman. The majority of Sunnī *fuqahā'* consider that prohibited (*ḥarām*), but the Shāfi'īs consider it permissible, reasoning (in part) that he would not fornicate with the other woman if she were available to him, and his act can be described as imagining that his wife has the appearance of the other woman. The Mālikī Ibn al-Ḥajj said this is a type of fornication. See: Ibn 'Ābidin, *Radd al-Muḥtār* (Cairo: Bābī Ḥalabī, 1966/1386), 6:372; Ibn Ḥajar al-Haytamī, *Tuḥfat al-Muḥtāj* (Cairo: al-Maktabah al-Tijāriyyah, 1983/1357), 7:205; Ibn al-Ḥājj, *al-Madkhal ilā tanmiyat al-a'māl bi-taṣīn al-niyāt* (Beirut: Dār al-Kutub al-'Ilmiyyah, 1995), 1:370. Although the majority do consider this *ḥarām*, nevertheless the sin would be expected to be less than that of literal fornication.

⁶⁰ Ibn 'Uthaymīn comments, "It is prohibited [for him] to end copulation before she climaxes, because he would be depriving her of full enjoyment." Ibn 'Uthaymīn, *al-Sharḥ al-Mumtī'* (Dammām: Dār Ibn al-Jawzī, 2002-7/1422-7), 12:417.

danger of addiction to masturbation, which is a known psychological disorder for which people are treated.

If a married person engages in a romantic, erotic or physical sexual relationship with an AI bot, then the sin is greater than if they were unmarried, because:

(1) they are abjuring the legitimate outlet for their desires and choosing a prohibited channel.

(2) they are probably neglecting the rights of the spouse (including the right to spend time together, and to receive emotional support).

وقال الجصاص: ومن المعروف أن يوفّيها حقّها من المهر والتّقنة والتّقسم، وترك أذاها بالكلام الغليظ، والإعراض عنها، والميل إلى

غيرها، وترك العيوس والقطوب في وجهها بغير ذنب

The Hanafi scholar al-Jaṣṣāṣ said, "Part of customary [good treatment of the wife] is that he fulfill her rights of *mahr*, financial support, and [fair] division [of time], and that he avoid hurting her with harsh words, abjuring / turning away from her, inclining to some other [woman], and frowning and scowling at her when she has not committed any sin."⁶¹

As for marriage to AI companions, this would be an invalid marriage contract, because

(1) The Quranic verse about spouses of our own kind (i.e. humans)

وَمِنْ آيَاتِهِ أَنْ خَلَقَ لَكُمْ مِنْ أَنْفُسِكُمْ أَزْوَاجًا لِتَسْكُنُوا إِلَيْهَا وَجَعَلَ بَيْنَكُمْ مَوَدَّةً وَرَحْمَةً إِنَّ فِي ذَلِكَ لَآيَاتٍ لِقَوْمٍ يَعْتَكَرُونَ

"And of His signs is that He created for you from yourselves mates that you may find tranquillity in them; and He placed between you affection and mercy. Indeed in that are signs for a people who give thought." Q.30:21

(2) The AI bot cannot give meaningful consent, because it is not really conscious.

Deepfakes and Digital Manipulation

A) Technical Background

AI technology has made it possible to not only generate realistic fake videos, but also manipulate existing videos to change faces or voices of people in the video. This can even be done in real-time, i.e. a livestream of a lecture can be manipulated instantly, so those watching online might be seeing or hearing something different to what is seen or heard by people who are physically on-site.

The most pervasive impact on daily life is the "Liar's Dividend." As the public becomes aware that any audio or video can be faked, the immediate defense for anyone caught in a compromising but authentic recording is to claim it is AI-generated. This creates a world where seeing is no longer believing. In daily life, this affects everything from HR disputes (claiming a recorded threat was a deepfake) to domestic litigation. The due diligence required of the average person is now strained.

The widespread accessibility of consumer-grade digital manipulation tools has led to considerable social harm. Examples include:

- Apps that allow users to generate sexually explicit imagery of acquaintances,

⁶¹ Jaṣṣāṣ, *Aḥkām al-Qur'ān* (Beirut: Dār al-Kutub al-'Ilmiyyah, 1994/1415), 2:138.

coworkers, or classmates using social media photos. The United States introduced the "Take It Down Act" (2025) and the DEFIANCE Act (2026) to create a federal, 48-hour takedown window for such content, with increased penalties for content involving minors. As of 2026, many of these apps are being actively removed from major app stores, but they continue to exist on the web and social media, often utilizing cryptocurrency for anonymous, untraceable transactions.

- Scammers can use audio clips from social media to clone the voice of a loved one. They then call a family member claiming to be in an accident or jail, asking for immediate funds. It can be incredibly difficult to trace these scammers.
- The ubiquitous use of real-time "beauty filters" that use generative adversarial networks (GANs) to alter facial and body structure in pictures and videos posted online - leading to a notable increase in dysmorphia and unrealistic expectations.

B) Fiqhī Dimensions

Regarding producing fake or deepfake pictures, video or other media for oneself, this is related to the ruling on pictures in general:

1. For those who prohibit all man-made pictures, it is haram.
2. For those who allow two-dimensional (flat) pictures,⁶² it would appear to be permissible as a form of entertainment, by analogy to exercising the imagination, as long as the content does not involve anything haram. E.g. It would not be permissible to use apps to undress photos of men and women (especially non-mahram), even if you don't share them with anyone else.
3. As for the view that prohibits all pictures, but makes an exception for photographs (by analogy to a reflection in a mirror), it depends on the details:
 - a. If AI is used to generate a picture (or video or audio) that is fake, but the people shown are real people, the ruling would appear to be the same in #2 above.
 - b. If AI is used to manipulate a picture (or video or audio) then it would be *haram*, because the deepfake is not a reflection of an actual person, and because the human creative element is involved.

The ruling on sharing fake or deepfake pictures, video or other media would be the same as the ruling on lying or spreading false information.⁶³ This is not permissible unless you make it clear to the people receiving it that this is AI-generated or AI-modified. A hadith gives 3 exceptions where lying is permissible.⁶⁴ Whether the use of AI-enhanced lying (in the form of AI-generated or AI-manipulated material) in these three cases gives cause for pause and reflection (موضع تأمل), but the authors of this paper are reluctant to

extend the permission.

As far as evidence in court cases, a key hadith is:

⁶² This is the view of the Mālikī *madhhab*, as well as some Ḥanbalis, and the Ja'fari *madhhab*.

⁶³ This is according to the view that flat pictures are permissible to produce, and also according to the view that flat pictures are permissible to use.

⁶⁴ For reconciling two people, in war, and small matters between a husband and wife. [Muslim]

الْبَيِّنَةُ عَلَى الْمُدَّعِي وَالْيَمِينُ عَلَى مَنْ أَنْكَرَ

"But the [burden of] proof (*bayyinah*) lies on the one who is making the claim, and the oath must be taken by the one who rejects [the claim]."⁶⁵

The majority of fuqaha' interpret *bayyina* (proof) in a narrow sense, as the testimony of two witnesses, because the Prophet once told al-Ash'ath ibn Qays, in the course of his dispute with another man:

شَاهِدَاكَ أَوْ يَمِينُهُ لَيْسَ لَكَ إِلَّا ذَلِكَ

"Your two witnesses, or his oath. You have no other [option]."⁶⁶

Hence, the norm in pre-modern Islamic law has been to privilege oral testimony over written and other forms of evidence. Written documents would typically have to be authenticated by oral testimony of the witnesses. With the increased availability of and importance given to documents in the modern period, especially the emergence of official notaries, there was a shift, and nowadays, officially notarized documents are admissible as evidence in the courts of many Muslim countries, although they may not be considered conclusive evidence, and can be challenged, due to the possibilities of coerced signing, forgery, etc. *Fuqahā* agreed that circumstantial evidence can be used as a weighting factor, but that most would not typically base a verdict solely on such evidence, but Ibn al-Qayyim seems to disagree.

A related issue, discussed by pre-modern *fuqaha'*, is that of whether the judge is allowed to issue a verdict based on his personal knowledge. Today, such knowledge could come from non-witness evidence. Imām Mālik said the judge is not allowed to issue a verdict based on his personal knowledge, and this is also one view from Imām Aḥmad. Imāms Abū Ḥanīfah and al-Shāfi'ī said the judge is allowed to rule based on his certain knowledge, provided the case does not involve one of the Rights of God (e.g. the *ḥudūd*). Hence, by agreement, video evidence would not be sufficient to convict someone of fornication or theft, or at least the suspect could not receive the prescribed punishment *ḥadd* (e.g. whipping). Nevertheless, if the video evidence was compelling, the judge might impose a lesser, discretionary punishment (*ta'zīr*) on the suspect. The rise of AI-generated videos, and deepfakes calls for more caution on the part of judges. Videos can be embedded with digital watermarks as a sign that they are genuine, and if a judge receives unbiased expert testimony that a video is almost certainly genuine, then he could potentially use it as evidence, but would not be allowed to impose a *ḥadd* punishment on its basis.

There are at least two other *fiqhī* issues related to this topic:

1. Caution requires that imams abstain from performing virtual weddings (over Zoom), because of the possibility that the Zoom stream is being deepfake manipulated in real time, to disguise the identity of one or more of the parties to the contract.
2. Among the general public: Awareness of fake and deepfake videos are a reminder to Muslims to not jump to conclusions if they see a compromising video of another Muslim, but rather to give them the benefit of the doubt. Imām Ja'far al-ṣādiq said,

إِذَا بَلَغَكَ عَنْ أَخِيكَ الشَّيْءَ تَنَكَّرْهُ، فَالْتَمَسْ لَهُ عَذْرًا وَاحِدًا إِلَى سَبْعِينَ عَذْرًا، فَإِنْ أَصْبَتْهُ، وَإِلَّا، قُلْ: لَعَلَّ لَهُ عَذْرًا لَا أَعْرِفُهَا

⁶⁵ Bayhaqī.

⁶⁶ Muslim.

"If objectionable [news] reaches you about your brother, then seek from one to seventy excuses for him that might apply to him. Otherwise, tell yourself: maybe he has an excuse [that I can't think of]."⁶⁷

Daily Conduct and Moral Harm

A) Technical Background

Whereas in the past, students could only cheat on assignments by copying, or having someone else do the work for them, now any student with internet access can provide a specific prompt to a LLM AI application (ChatGPT, Claude, etc) and within seconds generate a customized essay or slideshow. Subsequently, a number of websites and apps emerged to help teachers detect AI-generated content, but (1) these tools vary widely in their reliability, and (2) other apps were developed to humanize AI-generated output to make it more difficult to detect.

Dartmouth Provost Santiago Schnell recently observed that, "AI has not created new educational problems; it has made old ones impossible to ignore. The habit of rewarding performance over understanding, fluency over depth, and polish over genuine engagement was already present in our institutions before the first language model was trained. AI simply industrializes and accelerates those habits until their emptiness becomes undeniable."⁶⁸

Even some academics are using AI to help write journal articles, and this not only brings up concerns related to plagiarism and originality, but is also concerning because of the phenomenon of AI hallucinations (when AI invents false information, such as attributing information to non-existent books and articles). A recent study, that checked 111 million references across 2.5 million papers in academic journals, gave "a conservative estimate of 146,932 hallucinated citations in 2025 alone."⁶⁹ The authors comment that their "findings demonstrate that LLM hallucinations are infiltrating knowledge production at scale, threatening both the reliability and equity of future scientific discovery as human and AI systems draw on the existing literature."

AI can also be used to craft deceptive resumes, for example: by "upscaling" things a candidate has actually done but describing them in an inflated way, or generating fake employment histories, or fake work samples / portfolios. Because some companies use AI agents to sort or screen resumes received for job openings, some candidates stuff their resumes with keywords that will cause AI to rank their resume higher. Some candidates even insert into their resume invisible instructions (white text, or microscopic font) to the AI agents.

Direct use of generative AI to compose essays and the like is not the only form of AI

⁶⁷ Bayhaqī, *Shu'ab al-Īmān*. A similar statement is attributed to the Tabi'ī Muḥammad ibn Sīrīn by Ibn 'Asākir.

⁶⁸ Santiago Schnell, "Repairing the Ruins: Why AI Can't Replace Education," National Catholic Register, 4/13/2026. <https://www.ncregister.com/commentaries/schnell-repairing-the-ruins> . Accessed 6/12/2026.

⁶⁹ Zhenyue Zhao, Yihe Wang, Toby Stuart, Mathijs De Vaan, Paul Ginsparg, Yian Yin. "LLM hallucinations in the wild: Large-scale evidence from non-existent citations" 05/08/2026. arXiv:2605.07723 [cs.DL]. <https://doi.org/10.48550/arXiv.2605.07723> . Accessed 6/12/2026.

used by students. Spellcheckers and grammar checkers (such as Grammarly) often use AI technology.

AI is also widely used, especially by social media apps, to maximize user engagement with the app/platform. By attending to behavioral data at the millisecond level—such as watch time, clickthrough rate, or hover duration—AI models construct probabilistic models of specific user vulnerabilities to increase engagement. This creates a systemic moral hazard where the system incentivizes and facilitates addictive engagement patterns as a functional objective of the model. Because these models are optimized for retention rather than virtue, the black box remains indifferent to the ethical quality of the engagement. Whether the system is streamlining academic dishonesty via LLMs or surfacing haram content on an "Explore" page the user did not actively seek out, the underlying metric remains the same: the mathematical maximization of user persistence at the cost of informational or moral agency.

Engagement-driven algorithms, as well as conversations with AI chatbots, can reinforce a person's pre-existing biases, leading them to become more entrenched in a conspiracy theory that they had previously believed. However, honest conversations with AI (including using AI to help check facts) also have the potential to reduce belief in conspiracy theories.⁷⁰

Additionally, the weaponization of AI-driven bots has altered the landscape of the internet, and how users must engage with online information. These autonomous agents use LLMs and other techniques to conduct targeted phishing, simulate false public consensus, and scrape private data. By masquerading as human actors, these systems are used by malicious actors and erodes digital safety and trust.

B) Fiqhī Dimensions

I. Academic Dishonesty

"[L]earning is not the production of acceptable performances but the formation of a person capable of truth, judgment and responsibility."⁷¹

The statement above was made by a Catholic writer, but it resonates with Islamic values. It is common knowledge among Muslims that Islam values beneficial knowledge, not only religious but also worldly. The Prophet ﷺ was commanded (in what was, according to most ulema, the first Quranic revelation) to "Read!" Q.96:1, and the Prophet promoted literacy in the community. Not only is the goal of knowledge to produce an integral human, made whole by knowledge, but knowledge requires involvement of the whole person. As the Arabic tells us,

لَنْ يُعْطِيَكَ الْعِلْمُ بَعْضَهُ حَتَّى تُعْطِيَهُ كُلَّهُ

"Knowledge will not give you part of itself until you give it all of yourself."

This starts, ideally, with writing, as distinct from typing. Allah "taught by the pen" Q.96:4, and research has shown that people who take notes with pen and paper learn better than

⁷⁰ MIT Sloan Office of Communication, "Can AI talk us out of conspiracy theories?" 9/12/2024, <https://mitsloan.mit.edu/press/can-ai-talk-us-out-conspiracy-theories>, accessed 6/12/2026. Note the technical article in question, in the *Science* journal, is currently under review due to potential inconsistencies of data.

⁷¹ Schnell, *op. cit.*

those who take notes electronically.⁷² Writing by hand also engages neural circuits of the brain more than typing.⁷³

Use of AI makes people lazy, creating the illusion of knowledge, while sacrificing both depth and long-term retention. Even before AI, short-length videos (such as Facebook Reels, and TikTok) were gradually reducing people's attention-spans. A scientific study found that addiction to such videos impacts physical and mental health, and in particular it "could negatively impact self-control and diminish executive control".⁷⁴ With the now extensive use of AI, this intellectual laziness has taken on new dimensions. A scientific study divided people into three groups – (1) using LLM AI, (2) using only a conventional search engine, (3) human brain only – and asked them to write an essay. Those in the AI-use group showed lower brain connectivity, and also worse long-term performance (including memory).⁷⁵ Use of AI has a negative effect on critical thinking skills. If this is the effect on adults, who have already developed critical thinking skills, the danger for children is even more, because through regular use of AI to provide answers, they may never even develop the ability to think critically.

"Cognitive offloading" is a major danger of relying on AI to produce written material.⁷⁶ Meaningful learning needs us to meaningfully engage with the content: to think about different angles, to make comparisons and make connections to what we already know, and to think about the bigger picture.⁷⁷

Relying on LLMs cut out these steps, or actually offloads them to AI, thereby depriving us of the learning experience, and its joy. As it has been said, about that moment when something clicks in the brain,

فهمنا حين فهمنا

"We were ecstatic when we understood."⁷⁸

Cognitive offloading can reduce long-term memory, critical thinking, and cognitive endurance. Part of gratitude to Allah for giving us a functioning brain is to use it properly.

⁷² Horbury, S. R., & Edmonds, C. J. (2021). Taking Class Notes by Hand Compared to Typing: Effects on Children's Recall and Understanding. *Journal of Research in Childhood Education*, 35(1), 55–67. <https://doi.org/10.1080/02568543.2020.1781307>, accessed 6/12/2026.

⁷³ Marano, Giuseppe, Georgios D. Kotzalidis, Francesco Maria Lisci, Maria Benedetta Anesini, Sara Rossi, Sara Barbonetti, Andrea Cangini et al. "The neuroscience behind writing: Handwriting vs. typing—Who wins the battle?." *Life* 15, no. 3 (2025): 345, accessed 6/12/2026.

⁷⁴ Tingting Yan et. al. "Mobile phone short video use negatively impacts attention functions: an EEG study" in *Front Hum Neurosci*. 2024 Jun 27;18:1383913. doi: 10.3389/fnhum.2024.1383913. Accessed 6/24/2026.

⁷⁵ <https://www.media.mit.edu/publications/your-brain-on-chatgpt/> accessed 6/12/2026.

⁷⁶ For a non-technical article covering this, see: <https://www.bbc.com/future/article/20260505-how-to-use-ai-without-turning-your-brain-to-mush> . On the more technical side: <https://www.media.mit.edu/publications/your-brain-on-chatgpt/>, <https://policyoptions.irpp.org/2025/09/ai-memory/>, <https://www.media.mit.edu/publications/your-brain-on-chatgpt/>.

⁷⁷ Justin Sung, "How ChatGPT Slowly Destroys Your Brain", 8/2/2025, <https://www.youtube.com/watch?v=6sJ50Ybp44I>, accessed 6/9/2026.

⁷⁸ Hariri, *Maqāmāt* (Beirut: Dār al-Ma'ārif, 1873), 241.

"Oho they cried,
The world is wide,
But fettered limbs go lame!"⁷⁹

Just as muscles atrophy (degrade) when they are not used for an extended period of time, so too can the brain atrophy. Some people make the argument that freeing our brain of these tasks opens the door to moving on to different (possibly higher-level) tasks for our brain, such that we emerge better off. Whether that will be the case remains to be seen.

Someone who "learns" from an answer generated by AI might superficially and temporarily "know" it (*ilm*) but does not have encompassing knowledge (*ihāṭah*) of it. He is essentially practising blind following (*taqlid*),⁸⁰ and therefore does not really know because of the lack of confidence (*thiqah*).⁸¹

وَلَا يُنَبِّئُكَ مِثْلُ خَبِيرٍ

"None can inform you like an expert (*khābir*)" Q.35:14

Shaykh 'Abdul-Rahman al-Barrak explains, "The word *khābir* is more specific than *alīm* (knower), because it indicates knowledge of hidden aspects of matters, and their limits."⁸²

Also relevant here is the fact that AI systems are inherently biased, because the data they are trained on is human data, and humans have their own biases. This is especially important for religious matters, but also applies to non-religious biases.

وَإِنْ تُطِيعْ أَكْثَرَ مَنْ فِي الْأَرْضِ يُضِلُّوكَ عَنْ سَبِيلِ اللَّهِ إِنْ يَتَّبِعُونَ إِلَّا الظَّنَّ وَإِنْ هُمْ إِلَّا يَخْرُصُونَ

"And if you obey most of those upon the earth, they will mislead you from the way of Allah. They follow not except assumption, and they are not but falsifying." Q.6:116

Learning directly from humans makes it easier to detect bias, especially since many people tend to assume that AI is always presenting objective truth.

The above discussion makes a strong argument for not relying on AI at the expense of true learning, even if you are not submitting the work to a teacher or for publication. If you are actually submitting the work as your own, then it goes without saying that falls under the sin of deceit.

وَمَنْ تَحَلَّى بِمَا لَمْ يُعْطَ كَانَ كَلَّاسٍ ثَوْبِي زُورٍ

"Whoever pretends to have more than he actually has is like someone wearing two garments of lies."⁸³

Finally, it is worth pointing out the loss of originality (and subsequently loss of creativity). AI often generates unnatural prose – sometimes 'purple prose' with too many adjectives. It also forces diverse human individuals into a bland uniformity, which is a sort

⁷⁹ Oscar Wilde, "The Ballad of Reading Gaol."

⁸⁰ "People even trust AI over their own thinking and intuition – even when the AI is wrong – something researchers at the University of Pennsylvania in the US call "cognitive surrender". Thomas Germainn. "'Think outside the bots': How to stop AI from turning your brain to mush," BBC, 5/6/2026, <https://www.bbc.com/future/article/20260505-how-to-use-ai-without-turning-your-brain-to-mush>, last accessed 7/24/2026.

⁸¹ Askarī, *al-Furūq al-Lughawīyyah*: *alīm/muḥīṭ*, *ilm/taqlid*. See also: *ilm/tabyīn*

⁸² Barrāk, *Kitāb Tawqīh al-Masā'il al-'Aqdiyyah fī Muqaddimat al-Risālah al-Qayrawāniyyah* (Riyadh: Mu'assasat Waqf al-Shaykh al-Barrāk, 2021/1442), p.70.

⁸³ Abū Dāwūd.

of betrayal of our God-given individuality.

أَيُّهَا الْإِنْسَانُ مَا غَرَّكَ رَبِّكَ الْكَرِيمِ ۝ الَّذِي خَلَقَكَ فَسَوَّاكَ فَعَدَلَكَ ۝ فِي أَيِّ صُورَةٍ مَّا شَاءَ رَكَّبَكَ ۝

"O mankind, what has deceived you concerning your Lord, the Generous, (6) Who created you, proportioned you, and balanced you? (7) In whatever form He willed has He assembled you." Q.82:6-8

Total uniformity makes for a boring world, and unfortunately, as more people are using AI to generate emails, letters, articles, etc., these are starting to sound the same. Gone are the idiosyncrasies and unique styles of individuals. A further dimension is added to the problem when the AI-generated text published by humans is now used to train AI systems. There is the risk of AI becoming dumber, and humans losing even more creativity than we already have.⁸⁴

Nevertheless, it is not reasonable to expect that students and researchers will stop using AI entirely, but at least full disclosure should be made, i.e. if you use AI for an assignment or article, you should declare this openly, and also give details about how you used it. Acceptable uses of AI, that do not compromise academic honesty include:

- Using it to find the names of books, articles and other sources, although one should first make one's own effort, in order not to fall into cognitive offloading.
- Using spelling or grammar correction tools on text you have written yourself.

Regarding disclosing the use of AI, this is important to avoid deception (see the hadith quoted above) and also because we have to give thanks and credit where it is due. Although AI is not a person capable of being thanked, nevertheless we can consider that we are thanking the people who developed and trained the AI system.

من لم يشكر الناس لم يشكر الله

"Whoever does not thank people does not thank Allah."⁸⁵

We should use AI either

- Temporarily, e.g. use a spellchecker or grammar checker with attentiveness, learning from it, until we can wean ourselves off it.
- Intelligently, in a way that promotes good learning rather than cognitive offloading. Try to be a centaur (who has "structured and controlled interactions with AI, harnessing it as a tool for targeted efficiency"), not a cyborg (who contributes something but allows AI to play a leading role), and definitely not a self-automator (who offloads the work to AI). Centaurs learn more in the long run.⁸⁶

II. Use of AI for workplace-related issues

Using AI to clean up the structure of a resume would generally be permissible, analogous to using the services of a human editor or career consultant. However, use of AI

⁸⁴ Shumailov, Ilia, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. "The curse of recursion: Training on generated data makes models forget." arXiv preprint arXiv:2305.17493 (2023).

⁸⁵ Tirmidhī and others.

⁸⁶ For more details, see Dylan Walsh, "3 ways to use AI: Are you a cyborg, a centaur, or a self-automator?", 03/26/2026, <https://mitsloan.mit.edu/ideas-made-to-matter/3-ways-to-use-ai-are-you-a-cyborg-a-centaur-or-a-self-automator>, accessed 6/13/2026.

to 'improve' your resume in a dishonest way (e.g. upscaling, generating fake work histories or portfolios) is clearly prohibited, in light of what was mentioned above about deceit. The fact that other people are doing it does not mean it is now ok. As the fiqh maxim says,

لا عبرة بالعرف الفاسد

"An unsound customary norm is given no consideration."

And as the Sahabī Ḥudhayfah said,

لَا تَكُونُوا إِمَّعَةً، تَقُولُونَ: إِنَّ أَحْسَنَ النَّاسِ أَحْسَنًا، وَإِنْ ظَلَمُوا ظَلَمْنَا، وَلَكِنْ وَطَّنُوا أَنْفُسَكُمْ، إِنَّ أَحْسَنَ النَّاسِ أَنْ تُحْسِنُوا، وَإِنْ أَسَاءُوا فَلَا تَظْلِمُوا.

"Do not be a blind follower, who does good as long as other people are doing good, and is unjust when people are unjust. Rather, train yourselves to do good when people do good, but if [people] do evil, then do not commit injustice."⁸⁷

It should be noted that one's income would be *ḥarām* if one obtained a job through false pretenses or deception.

Use of bots to make posts on social media, for propaganda purposes or to sway public opinion, would be prohibited if the content of the posts is false or deceptive. Even if the content of the posts is true, there is still an element of deception, because one is giving the impression that this comment is being made by an actual human being. If the user name and profile makes it clear that this is a bot, and the content of the post is not false nor misleading nor haram, then there does not seem to be grounds for blanket prohibition. Rather, the ruling would depend on the intentions and objectives. Similarly, it is prohibited to use bots for the purposes of defrauding or stealing private information.

Use of AI for forging legal documents is similarly prohibited, and probably even more serious, because it (typically) would involve depriving another person of his rights.

مَنْ اقْتَطَعَ حَقَّ امْرِئٍ مُسْلِمٍ بِيَمِينِهِ فَقَدْ أَوجَبَ اللَّهُ لَهُ النَّارَ وَحَرَّمَ عَلَيْهِ الْجَنَّةَ. فَقَالَ لَهُ رَجُلٌ: وَإِنْ كَانَ شَيْئًا يَسِيرًا يَا رَسُولَ اللَّهِ قَالَ: "وَإِنْ قَضَيْتُ مِنْ أَرَاكِ"

"If anyone seized - by his oath - what rightly belongs to a Muslim, Allah has made the Hell-fire compulsory for him and forbidden for him the Paradise." A man asked, "O Allah's Messenger, even if it were something insignificant?" He replied, "Even if it were a twig from a mustard (*arak*) tree."⁸⁸

"مَنْ حَلَفَ عَلَى يَمِينٍ يَقْتَضِي بِهَا مَالَ امْرِئٍ مُسْلِمٍ هُوَ فِيهَا فَاجِرٌ لَقِيَ اللَّهَ وَهُوَ عَلَيْهِ غَضَبَانٌ"

"If anyone swears a firm oath - knowingly, intentionally, taking by it property belonging to a Muslim, Allah will be angry with him when he meets Him (on the Day of Resurrection)."⁸⁹

Workers who are required by their employers to use AI may do so even if they would prefer not to, as this would fall under the classification of widespread need or necessity. A

⁸⁷ Tirmidhī narrated it as a hadith, but most hadith authorities say the transmission to the Prophet is weak.

⁸⁸ Muslim.

⁸⁹ Bukharī, Muslim.

plausibility argument is provided by the fiqh maxim:

الحاجة تُنزل منزلة الضرورة، عامة كانت أو خاصة

“A need might take the same ruling as necessity if it is widespread.”

III. Responsibility

If the app or website can be used for honest and legal things, then the user who uses it for dishonest / illegal purposes bears sole responsibility. However, if the developers of the app/website know that the primary use will be for deceptive or illegal activities, then they are partially responsible. Ideally, there should be regulations banning such AI systems.

III. Algorithmic Exposure

There is a real danger of becoming intellectually and psychologically isolated in the echo chamber that AI produces from the person's own biases, whether it takes the form of conspiracy theories, or any other mistaken beliefs such as polytheism or terrorism. This frightening aspect of AI adds a new dimension to the Islamic theological dimension of *istidraj*, one definition of which is:

استجلاب الإنسان ما يشتهيهِ لايقاعه في الهلاك غير المتطَر

“Procurement of what the person desires, in order that he falls into unexpected ruin.”⁹⁰ This concept can be found in number of verses in the Quran, e.g.:

أَفَمَنْ رُئِيَ لَهُ سُوءُ عَمَلٍ فَرَأَاهُ حَسَنًا فَإِنَّ اللَّهَ يُضِلُّ مَنْ يَشَاءُ وَيَهْدِي مَنْ يَشَاءُ

“Then is one to whom the evil of his deed has been made attractive so he considers it good [like one rightly guided]? For indeed, Allah sends astray whom He wills and guides whom He wills.” Q.35:8

مَنْ يُضِلِّ اللَّهُ فَلَا هَادِيَ لَهُ وَيَذَرُهُمْ فِي طُغْيَانِهِمْ يَعْمَهُونَ

“Whoever Allah sends astray - there is no guide for him. And He leaves them in their transgression, wandering blindly.” Q.7:186

وَقَيَضْنَا لَهُمْ قُرْنَاءَ فَزَيَّنُوا لَهُمْ مَا بَيْنَ أَيْدِيهِمْ وَمَا خَلْفَهُمْ وَحَقَّ عَلَيْهِمُ الْقَوْلُ فِي أُمَمٍ

قَدْ خَلَتْ مِنْ قَبْلِهِمْ مِنَ الْجِنِّ وَالْإِنْسِ إِنَّهُمْ كَانُوا خَاسِرِينَ

“And We appointed for them companions who made attractive to them what was before them and what was behind them [of sin], and the word has come into effect upon them among nations which had passed on before them of jinn and men. Indeed, they [all] were losers.” Q.41:25

Indeed, we might be approaching the time mentioned in the hadith,

حتى إِذَا رَأَيْتُ شُحًّا مُطَاعًا، وَهَوًى مُتَّبَعًا، وَدُنْيَا مُؤْتَرَةً، وَإِعْجَابَ كُلِّ ذِي رَأْيٍ بِرَأْيِهِ، فَعَلَيْكَ بِخَاصَةِ نَفْسِكَ وَدَعِ الْعَوَامَّ

⁹⁰ *Mawsu'at al-Muṣṭalaḥāt al-Islāmiyyah* (Riyadh: al-Maktab al-Ta'āwunī lil-Da'wah wa-l-Irshād, 1441), p. 199.

"...until you see avarice obeyed, desires followed, the world preferred [over religion], and everyone is amazed with his own view. Then you should be worried about yourself in particular, and don't worry about the public."⁹¹

The antidote to this is multi-pronged:

- Strive to obtain authentic knowledge, and ponder over what you learn self-critically, so you are not a blind follower in fundamental matters of belief.⁹²
- Establish a regular routine of worship to give you spiritual strength and self-awareness.

إِنَّ الصَّلَاةَ تَنْهَى عَنِ الْفَحْشَاءِ وَالْمُنْكَرِ

"Indeed the prayer restrains from indecent and wrongful conduct." Q.29:45

Ibn Taymiyyah used to spend a lot of time on the morning *adhkār* (remembrances and supplications) and considered it his spiritual nourishment. A spiritually healthy and sound mind perceives reality more clearly. When people asked Ahmad ibn Hanbal, that after he passed away, to whom should they take their questions, he replied, "Ask 'Abdul-Wahhāb [al-Warrāq], for he is a pious man, and people like him are granted *tawfīq* (ability, success) to attain the truth."⁹³

- Root yourself in a religious community that is on the right path, at least in major issues.

عليك بالجماعة فإنما يأكل الذئب من الغنم القاصية

"You should stick to the *jamā'ah*, for the wolf devours the lone sheep."⁹⁴
Surrounding yourself with good people and interacting with them helps keep you in check, especially if the community is performing the duty of enjoining good and forbidding evil. Sincere friends can also help challenge your thinking. As the proverb says,

صديقك من صدقك لا من سددك

"Your friend is the one who tells you the truth, not the one who [always] affirms you."

- Regularly pray sincerely for guidance. Two relevant supplications are:
 - The Prophetic supplication:

اللَّهُمَّ رَبَّ جِبْرَائِيلَ وَمِيكَائِيلَ وَإِسْرَافِيلَ فَاطِرَ السَّمَوَاتِ وَالْأَرْضِ عَالِمَ الْغَيْبِ وَالشَّهَادَةِ أَنْتَ تَحْكُمُ بَيْنَ عِبَادِكَ فِيمَا كَانُوا فِيهِ يَخْتَلِفُونَ اهْدِنِي لِمَا اخْتَلَفَ فِيهِ مِنَ الْحَقِّ بِإِذْنِكَ إِنَّكَ تَهْدِي مَنْ تَشَاءُ إِلَى صِرَاطٍ مُسْتَقِيمٍ

"O God, Lord of Gabriel, and Michael, and Raphael, the Originator of the heavens and the earth, Knower of the unseen and the seen! It is You who judge between Your servants in that about which they disagree. Guide me, to that of the truth concerning

⁹¹ Tirmidhī (*ḥasan*), Ibn Hibban, and others. There is disagreement over the authenticity.

⁹² See the discussion on *ujb* (conceit) in: Maqdisī, *Mukhtaṣar Minhāj al-Qasidīn*. Ghazālī, *Ihyā 'Ulūm al-Dīn*.

⁹³ Ibn al-Jawzī, *Manāqib al-Imām Ahmad* (Dār Ibn Khaldūn), p. 438.

⁹⁴ Ahmad. Graded as *ḥasan* by al-Nawawī, *ḥasan li-ghayrih* by al-Arnā'ūt, *ṣaḥīḥ* by Ibn al-Qayyim, but *ḍa'īf* by al-Albānī.

which there has been disagreement, with Your permission, for verily You guide whom you will to a Straight Path.”⁹⁵

- The famous non-prophetic supplication:

اللهم أرني الحق حقاً وارزقني إتباعه ، وأرني الباطل باطلاً وارزقني اجتنابه

O Allah! Show me the truth as truth and grant me to follow it. Show me the falsehood as falsehood and grant me to avoid it.”

Shaykh Yasir Qadhi, in one of his khutbahs, mentioned a technique to use the algorithms to one’s religious benefit: make sure to often visit websites with good Islamic content (e.g. Quran recitations) so that the algorithm gives you more suggestions like this. And of course if you visit pornographic or Islamophobic sites often, you have yourself to blame when the algorithm feeds you more haram content.

⁹⁵ Muslim.

Conclusion

This paper has explored some common issues involving the use of AI by the general public in daily life, as opposed to more specialized types of AI. The selection of issues was necessarily limited, for reasons of practical necessity. One issue not covered in the paper, but which is equally important and deserves further study is the fiqh of privacy and online identity in the age of AI.

As mentioned earlier, the AI field is in a state of rapid flux, and readers are cautioned that the fiqh discussions in the paper are based on the current state of the field, and might become inapplicable if the technology changes significantly. In this paper, we have proposed some suggestions for consideration as the field rapidly develops, and encourage multi-disciplinary efforts to study the effects of AI on daily life and broader society.

We call on tech leaders, scholars, and the general public to be intentional and cautious in their development and use of AI technologies, and not to jump on the bandwagon and do things simply because other people are doing them. We also remind Muslims that we should not let our use of AI eclipse or compromise (1) our human connections, especially family relationships, (2) our mental health, (3) our mental activity and critical thinking skills. We should also consider how our use of AI impacts our ummah and the environment. And above all, our use of AI should not degrade our connection to Allah.