



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

The Fiqh of AI in Daily Use

Dr. Suleiman Hani
Dr. Mohammed Ansari

Dr. Suleiman Hani is the Dean of Academic Affairs at AlMaghrib Institute and a resident scholar in Michigan. He began studying the Islamic sciences from a young age and holds traditional ijāzāt in various disciplines, complemented by graduate degrees from Harvard University, the University of Jordan's College of Shari'ah, and a doctorate in leadership from NYU. His research and lectures, delivered across more than fifty countries, focus on navigating contemporary issues, intellectual and political thought, and ethical leadership, and he has authored numerous books and articles in these fields.

Dr. Mohammed Ansari holds a PhD in computer sciences from the University of Wisconsin-Madison with a specialization in artificial intelligence. He holds a senior scientist role in the industry and is passionate about developing technological solutions that rise to address contemporary challenges for the umma.

"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

Abstract	5
1. Introduction: Why Daily-Use AI Is an Urgent Fiqh Issue.....	6
2. Definitions and Scope: AI, Daily-Use AI, and Non-Moral Agency	9
A Technical Note on Forms of AI Relevant to This Paper	11
3. Methodological Foundations for Iftā' on Daily-Use AI	13
3.1 The Methodological Foundation	13
3.3 A Note on Governing Maxims for Daily-Use AI	19
4. AI as Moral Environment	21
4.1 Attention	21
4.2 Truthfulness and Epistemic Discipline	22
4.3 Privacy and Exposure	22
4.4 Social Relations	23
4.5 Ḥayā' and Relational Boundaries	23
4.6 Self-Presentation and Identity	24
4.7 Dependency and Diminished Agency	24
4.8 Scale.....	24
5. AI Virtual Relationships	25
5.1 Classification	25
5.2 AI Companionship.....	26
5.3 Romantic AI Usage.....	28
5.4 Minors and Vulnerable Users	30
5.5 AI Pets	31
5.6 Mental Health	31
5.7 Restoration of Family Photos	32
5.8 Practical Fiqh Framework	33
6. Deepfakes and Digital Manipulation	34
6.1 Truthfulness.....	34
6.2 Deception and Fraud	35
6.3 Sexual Deepfakes	35
6.4 Deepfakes and Epistemology	36
6.5 Guidelines.....	37
6.6 Religious Deepfakes	38
6.7 Consent Over One's Likeness.....	38
6.8 Modern governance frameworks	40
6.9 Psychological Impact of Deepfake Abuse.....	41
6.10 Civic Engagement	41
6.11 Practical fiqh framework.....	42
7. Daily AI Conduct and Moral Harm.....	44
7.1 Deception of Competence	44
7.2 Deception in Professional, Commercial, and Legal Communication	46
7.3 Impersonation and False Relational Communication	47
7.4 Algorithmic Exposure to Prohibited or Harmful Content	47

7.5 Addictive or Compulsive Engagement	48
7.6 Negligence Through Overreliance	49
7.7 The Spread of Harmful AI-Generated Content.....	49
7.8 Individual Responsibility and Platform Responsibility	49
7.9 The Principle of <i>Taklif</i>	50
7.10 Practical fiqh framework.....	50
8. Practical Guidelines for Daily-Use AI.....	52
8.1 Guidelines for Individuals.....	52
8.2 Guidelines for Families.....	53
8.3 Guidelines for Religious Leadership	53
8.4 Guidelines for Educational Institutions	54
8.5 Guidelines for Islamic Institutions	55
8.6 Guidelines for Developers and Technologists	56
9. Conclusion.....	58
Limitations of This Paper	59
Bibliography	60

Abstract

Artificial intelligence has moved from specialized technical settings into the ordinary moral environment of Muslim daily life. Generative, conversational, recommender, and synthetic-media systems now shape how people write, study, communicate, seek companionship, verify claims, present themselves, and encounter temptation. This paper develops a ten-step method of *iftā'* for evaluating daily-use AI, proceeding from *taṣwīr al-mas'alah* and *tahqīq al-manāṭ* through *fiqh al-wāqī'*, the baseline ruling, the *taklīfī* and *waḍ'ī* layers, *maqāṣid al-Sharī'ah*, *qawā'id fiqhīyyah*, *i'tibār al-ma'ālāt* and *sadd al-dharā'i'*, and *fiqh al-awlawiyyāt* and *al-muwāzanāt*, to *tanzīl al-ḥukm*. It argues that AI possesses no moral agency, *niyyah*, *dhimmah*, or legal capacity; responsibility remains with accountable human users, developers, platforms, institutions, and communities.

The method is applied, in sustained conversation with current technical and empirical research, to four domains: 1) the general moral ecology of daily-use AI; 2) emotional and intimate relationships with AI companions, evaluated against the classical regulation of desire and the precursors of *zinā*; 3) deepfakes and synthetic media, including an analysis of sexualized deepfakes under *qadhf*, *ta'rīd*, and fabricated evidence, and the protection of religious transmission against false attribution; and 4) AI-facilitated deception, academic dishonesty, impersonation, addictive engagement, and harmful exposure, including the *tabayyun* owed to those accused of misconduct by unreliable AI detectors.

Each domain concludes with a practical fiqh framework, and the paper closes with operational guidelines for individuals, families, imams, educational and Islamic institutions, and Muslim technologists. The paper rejects both technological panic and uncritical adoption: daily-use AI may be permissible when it assists lawful purposes without deception, harm, dependency, privacy violation, or misrepresentation; blameworthy when it weakens discipline or healthy human relationships; and prohibited when it becomes a means to falsehood, unlawful intimacy, exploitation, injury, fraud, religious misguidance, or serious foreseeable harm.

Disclaimer: AI-assisted tools were used for organization, literature discovery, proofreading, and citation-format support. The authors retained responsibility for all writing, analysis, source verification, and legal-ethical conclusions.

1. Introduction: Why Daily-Use AI Is an Urgent Fiqh Issue

Artificial intelligence has moved with unusual speed from specialized technical domains into the ordinary texture of daily life. It now assists users in writing, searching, translating, designing, studying, planning, soliciting advice, socializing, entertaining themselves, and making decisions. Generative AI systems in particular have intensified this transition by allowing ordinary users to produce fluent text, images, audio, video, code, simulated personalities, and synthetic forms of interaction with little technical expertise. The result is not simply a new set of tools, but a new layer of mediation between the human person and knowledge, intimacy, truth, self-presentation, privacy, and moral agency. The 2025 Stanford AI Index noted that organizational AI use rose sharply in 2024, while generative AI investment and deployment continued to expand across sectors.¹ This acceleration makes AI no longer a distant policy question for governments and corporations alone. It has become a practical question for families, students, educators, imams, counselors, employers, and ordinary Muslims seeking to live with integrity in a rapidly altered moral environment.

The urgency of this paper arises from the fact that daily-use AI does not merely help users perform isolated tasks. It can shape attention, lower barriers to deception, simulate companionship, imitate authority, produce persuasive falsehood, intensify dependency, and blur the boundary between assistance and substitution. In many cases, the relevant fiqh question is therefore not only whether a particular output is permissible, but what form of conduct the system habituates, what relationships it displaces, what rights it violates, what harms it makes easier, and what long-term consequences it produces when used repeatedly and at scale. Islamic jurisprudence has always been concerned not merely with the external form of acts, but with their causes, means, effects, rights, harms, and consequences. This makes the question of daily-use AI properly a matter of fiqh, akhlāq, and iftā', not merely a matter of technical convenience.

The Sharī'ah does not approach new technologies through uncritical rejection or uncritical acceptance. The default treatment of worldly tools and customary means is not suspicion for its own sake. At the same time, permissibility in origin does not immunize a tool from restriction when it becomes a means to deception, unlawful desire, privacy violation, reputational injury, intellectual corruption, exploitation, or foreseeable harm. The classical tradition provides a disciplined vocabulary for this work: *taḥqīq al-manāt*, *fiqh al-wāqī'*, *sadd al-dharā'ī'*, *i'tibār al-ma'ālāt*, *maqāṣid al-Sharī'ah*, *qawā'id fiqhiyyah*, *amānah*, *ḍamān*, and the distinction between direct and indirect causation. These concepts prevent the discussion from collapsing into either technological panic or permissive naiveté.

In this respect, AI resembles but also exceeds earlier digital technologies. Social media already demonstrated that technologies marketed as neutral tools of connection can restructure attention, public speech, reputation, community formation, and moral habits. AI systems inherit many of these concerns, but add new capacities: generative output, synthetic identity, conversational simulation, (a form of) reasoning, memory across conversations, personalization, automated persuasion, and increasingly autonomous, goal-

¹ Stanford Institute for Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2025* (Stanford, CA: Stanford University, 2025). The report notes sharp growth in organizational AI use and continued expansion of generative AI investment.

driven behavior. The ethical challenge is therefore not simply that users may employ AI for prohibited ends. It is that AI systems may be designed to anticipate desire, prolong engagement, imitate intimacy, automate misrepresentation, and produce plausible but unreliable information. Langdon Winner's well-known argument that artifacts can embody political and social arrangements is especially relevant here, as is later work on persuasive technology showing how digital systems can be deliberately designed to shape user behavior.²

For Muslim jurists, the matter is weightier still because the Shari'ah evaluates human action within a dual horizon: worldly welfare and ultimate accountability before Allah. Ibn al-Qayyim famously described the Shari'ah as founded upon wisdom and the welfare of servants in this life and the next: it is entirely justice, mercy, welfare, and wisdom. Thus, any ruling or application that departs from justice into oppression, mercy into its opposite, welfare into harm, and wisdom into futility is not truly from the Shari'ah, even if it is inserted into it through interpretation.³ This statement is not a license for subjective moralism, but a reminder that fiqh must remain faithful to both revealed texts and the purposes those texts secure. Daily-use AI is precisely the kind of phenomenon that requires such disciplined fidelity: the texts provide the red lines and moral anchors, while the jurist must also understand the reality of the technology and the consequences of its use.

This paper therefore examines four interrelated areas of daily-use AI: first, the general scope of AI's impact on personal conduct, social relations, privacy, and harm; second, emotional or intimate relationships with AI companions and virtual partners; third, deepfakes and synthetic media that undermine truth, testimony, reputation, and trust; and fourth, AI-facilitated deception, academic dishonesty, impersonation, addiction-like engagement, and exposure to harmful or prohibited content. These topics are treated together because they share a deeper moral structure. Each involves a technology that can benefit users in some contexts while also reshaping conduct, weakening restraint, obscuring responsibility, or harming others in ways that are not always immediately visible.

The purpose of the paper is not to declare sweeping rulings detached from technical reality, nor to outsource Islamic judgment to secular AI ethics. Rather, it seeks to offer a Sunni juristic framework capable of issuing nuanced guidance while remaining rooted in the sources and methods of the tradition. Contemporary governance documents such as the OECD AI Principles, the NIST AI Risk Management Framework, and UNESCO's Recommendation on the Ethics of Artificial Intelligence are useful in identifying technical and social risks, especially around transparency, human oversight, privacy, dignity, and risk management.⁴ Yet these frameworks cannot replace the Shari'ah's account of worship, moral responsibility, chastity, truthfulness, amanah, family integrity, and accountability before Allah. They may help describe aspects of the *wāqi'*, but they do not determine the

² Langdon Winner, "Do Artifacts Have Politics?," *Daedalus* 109, no. 1 (Winter 1980): 121–36; B. J. Fogg, *Persuasive Technology: Using Computers to Change What We Think and Do* (San Francisco: Morgan Kaufmann, 2003).

³ Ibn al-Qayyim al-Jawziyyah, *I'lām al-Muwaqqi'īn 'an Rabb al-'Ālamīn*, ed. Muḥammad 'Abd al-Salām Ibrāhīm, 4 vols. (Beirut: Dār al-Kutub al-'Ilmiyyah, 1991), 3:11.

⁴ Organisation for Economic Co-operation and Development, *Recommendation of the Council on Artificial Intelligence*, OECD/LEGAL/0449, adopted May 22, 2019, amended May 3, 2024; National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1 (Gaithersburg, MD: NIST, January 2023); UNESCO, *Recommendation on the Ethics of Artificial Intelligence* (Paris: UNESCO, 2021).

ḥukm.

The central claim of this paper is that daily-use AI should be evaluated as a morally consequential technology whose ruling depends not only on the user's immediate intention, but also on the nature of the use case, the rights implicated, the harms facilitated, the system's design, the probability and severity of consequences, and the availability of safer alternatives. A daily-use AI application may be permissible when it assists a lawful purpose without deception, harm, dependency, privacy violation, or unlawful exposure. It may be disliked when it weakens discipline, replaces healthier human relationships, or habituates morally corrosive patterns. It may be prohibited when it enables falsehood, impersonation, sexualized interaction, betrayal of trust, academic fraud, violation of privacy, or foreseeable harm to individuals or the public. Such gradation is not indecision. It is the ordinary work of *iftā'* when a single technology contains multiple possible uses, contexts, and consequences.

2. Definitions and Scope: AI, Daily-Use AI, and Non-Moral Agency

For the purposes of this paper, artificial intelligence refers to machine-based systems that, for explicit or implicit objectives, infer from the inputs they receive how to generate outputs such as predictions, content, recommendations, or decisions that may influence physical or virtual environments. This definition aligns with the updated OECD formulation, which has become influential in contemporary AI governance.⁵ In practical terms, the paper is concerned with AI systems as ordinary users encounter them in daily life, including generative, conversational, recommender, companion, agentic, and synthetic-media systems. These forms are distinguished more precisely in the Technical Note below. Such systems may appear conversational, adaptive, and even emotionally responsive, but their apparent fluency must not be confused with consciousness, moral perception, spiritual insight, or legal accountability.

This distinction is indispensable for Islamic legal analysis. In contemporary technical literature, the word “agent” may describe a system that performs tasks, pursues programmed objectives, or acts semi-autonomously within an environment.⁶ In fiqh and uṣūl, however, moral and legal responsibility attach to the *mukallaf*: the accountable human subject addressed by divine command, possessing the relevant capacity for understanding, intention, and action. An AI system, no matter how fluent or autonomous in the technical sense, does not possess *niyyah*, *taqwā*, *dhimmah*, *wilāyah*, *shahādah*, or *ahliyyah*. It is not rewarded, punished, sinful, pious, just, corrupt, or morally answerable before Allah. The Qur’anic principle that Allah does not burden a soul beyond its capacity speaks to accountable beings, not artifacts or machines.⁷

Accordingly, when this paper discusses “AI harm,” “AI deception,” or “AI companionship,” the language should be understood as shorthand. The *ḥukm* does not attach to the machine as a morally responsible subject. It attaches to the human beings and institutions who design, deploy, market, rely upon, misuse, recommend, or knowingly benefit from the system. The relevant actors may include the individual user, the developer, the platform, the company, the educator, the employer, the imam, the parent, the counselor, or the institution that integrates the tool into its work. This distinction prevents two opposite errors: anthropomorphizing AI as though it were morally accountable, and absolving humans because the immediate output was generated by a machine.

The phrase “daily-use AI” in this paper refers to AI systems encountered in ordinary personal, educational, social, family, religious, and professional life. This includes AI writing assistants, chatbots, personal assistants, recommendation systems, image and video generators, voice cloning tools, translation systems, tutoring applications, productivity tools, companion bots, virtual characters, content filters, social media algorithms, and synthetic media tools. It also includes cases where AI is embedded invisibly into familiar platforms, such as search engines, messaging applications, social feeds, cameras, shopping tools, workplace software, and educational platforms. The concern is not only with intentionally

⁵ OECD, *Recommendation of the Council on Artificial Intelligence*, OECD/LEGAL/0449.

⁶ Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. (Hoboken, NJ: Pearson, 2021), chap. 2.

⁷ Qur’an 2:286. For the uṣūlī discussion of the legally accountable subject and legal capacity, see Wahbah al-Zuhayli, *Uṣūl al-Fiqh al-Islāmī*, 2 vols. (Damascus: Dar al-Fikr, 1986).

“using AI,” but with living in an environment increasingly mediated by AI.

This scope excludes, except by analogy, several important but distinct domains: autonomous weapons, state surveillance systems, judicial AI, medical triage, autonomous vehicles, AI-assisted *iftā'*, AI-generated fatwas, and large-scale institutional AI governance. Those topics require their own treatment because they involve different levels of risk, authority, expertise, and liability. Nevertheless, the boundary is not absolute. A daily-use AI tool may become high-stakes when used to generate evidence, produce religious claims, imitate a scholar, advise a vulnerable person, influence marital conduct, assist academic certification, or create reputational harm. The classification therefore depends on the use case, not merely the commercial label of the tool.

The paper also distinguishes between AI as a tool, AI as a medium, and AI as an environment. As a tool, AI may assist a user in a specific lawful task, such as summarizing notes, improving grammar, translating a permissible text, or organizing a schedule. As a medium, AI may shape communication between people, including how they present themselves, verify claims, or form impressions of others. As an environment, AI may influence the user's habits, desires, attention, dependencies, expectations, and moral thresholds. This third level is especially important. Some harms are not evident in a single prompt or response, but become clear only through repetition, personalization, escalation, and dependency.

This is why a sound fiqh treatment of daily-use AI cannot stop at the question: “Is the content generated by AI halal or haram?” That question is necessary, but incomplete. The jurist must also ask: Was there deception? Was a right violated? Was a person impersonated? Was consent obtained? Was privacy breached? Was a vulnerable user exploited? Was a prohibited desire stimulated? Was trust misplaced in a system known to hallucinate or fabricate? Was a student, employer, spouse, judge, imam, or audience misled about the source or reliability of the output? Was the use occasional and bounded, or habitual and dependency-forming? Was the system designed to assist the user's lawful agency, or to capture attention and manipulate behavior? Did the user of the system ascribe to its output greater reliability than claimed?

The distinction between assistance and substitution is also central. AI assistance may be permissible where it supports a lawful human act without misrepresentation. AI substitution becomes problematic where the system replaces a human responsibility that must be performed by the accountable person. A grammar tool that improves clarity is not the same as a student submitting AI-generated work as proof of personal competence. A scheduling assistant is not the same as an AI companion replacing marital intimacy. A labeled synthetic reenactment is not the same as a deepfake designed to deceive. A mental health support prompt is not the same as presenting a chatbot as a therapist, mufti, spouse, or spiritual guide.

This distinction also affects the ruling on reliance. AI output may be used as a preliminary aid in low-stakes contexts, but reliance becomes more restricted as the stakes rise. This is consistent with the broader risk-based logic recognized even in secular AI governance frameworks, where the consequences of error, opacity, bias, or misuse determine the level of required oversight. NIST's AI Risk Management Framework treats AI risk as affecting individuals, organizations, and society, and its generative AI profile specifically addresses risks heightened by generative systems, including harmful content,

information integrity, privacy, and misuse.⁸ From a fiqh perspective, this corresponds to the need to distinguish between low-risk assistance and high-risk delegation where rights, religion, reputation, family, property, safety, or public trust may be affected.

This layered approach protects the paper from two common weaknesses in contemporary discussions of AI: treating new technologies as though they automatically escape inherited categories, and forcing them into inherited categories without proper *taṣwīr*. The first error produces novelty without roots. The second produces rulings without reality. A serious Sunni treatment must do neither. It must begin by understanding the technology, identifying the legally relevant attributes, locating the closest juristic categories, weighing benefits and harms, considering consequences, and then issuing guidance with the appropriate level of confidence, limitation, and humility.

A Technical Note on Forms of AI Relevant to This Paper

Although the term “artificial intelligence” is often used as if it names a single technology, the daily-use cases examined in this paper involve several overlapping systems whose legal attributes differ. Generative AI refers to systems that produce text, images, audio, video, code, or other content in response to prompts or environmental inputs. Large language models are a major form of generative AI trained to predict and generate token sequences, most commonly text; multimodal foundation models extend related generative and interpretive capacities across text, image, audio, video, and other data types. In other words, these models produce well-received outputs – which in the past was a stand-in for competence – without necessarily having authentic content.

Recommender systems do not merely answer user prompts; they select, rank, and personalize what users encounter, often shaping attention and behavior. Companion bots and virtual partners combine conversational AI with memory, affective language, personalization, and sometimes romantic or erotic design, making their moral significance different from ordinary productivity tools. Agentic AI refers to systems that can pursue goals across steps, use tools, call external services, or act semi-autonomously within a digital environment. In technical literature, such systems may be called “agents,” but this is a technical use of the term and does not alter the legal conclusion that AI is not a *mukallaf*.

These distinctions matter because the ruling follows the legally relevant attributes of the use case, not the marketing label of the technology. A writing assistant, a recommender feed, a romantic chatbot, a voice-cloning tool, a synthetic-video generator, and an autonomous task agent do not raise identical fiqh questions. Some primarily implicate authorship and competence; others implicate privacy, attention, desire, emotional dependency, false attribution, evidentiary trust, or third-party harm. Likewise, not all AI risk arises from the same mechanism. Some risks come from probabilistic generation and hallucination, others from biased training data, insecure deployment, persuasive design, retrieval failures, anthropomorphic presentation, memory features, lack of provenance, or

⁸ National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (Gaithersburg, MD: NIST, January 2023); National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1 (Gaithersburg, MD: NIST, July 2024).

business models that reward engagement over wellbeing. A sound fiqh analysis must therefore avoid treating “AI” as a monolith. It must identify the actual system, function, user, context, rights affected, level of reliance, and foreseeable consequences before assigning the ruling.

3. Methodological Foundations for Iftā' on Daily-Use AI

A credible fiqh treatment of daily-use artificial intelligence must begin with method before ruling. The jurist who moves directly from a new technology to a legal conclusion without first establishing the nature of the act, the reality of the technology, the relevant causes, the affected rights, and the likely consequences risks issuing guidance that is either too permissive, too restrictive, or misdirected. This is especially true in the case of AI, where the same underlying tool may be used for lawful assistance, deceptive misrepresentation, emotional dependency, privacy violation, religious confusion, or reputational harm.

3.1 The Methodological Foundation

The task of *iftā'* in this domain therefore requires a layered method that combines fidelity to the sources with disciplined attention to reality. Although it is established in *iftā'* training and commonly known, one must remember that the following methodology is not sequential, it is an iterative and mutually informing process rather than a linear one. Furthermore, not all ten steps below will appear in the juristic process in an apparent manner, nor to the same degree of emphasis, depending on the issue being studied.

3.1.1 Taṣwīr Al-Mas'alah

The first step is *taṣwīr al-mas'alah*, the accurate conceptualization of the issue. Classical jurists repeatedly emphasized that sound judgment depends on first understanding the reality of the matter judged. A ruling on "AI use" in the abstract is rarely meaningful, because AI is not a single act. The relevant question may involve a student submitting AI-generated work, a married person forming a romantic relationship with an AI companion, a person creating a synthetic image of another without consent, an imam relying on an AI summary, a parent using AI to monitor a child, or a company designing a recommender system that drives users toward addictive, prohibited, or biased content. Each case carries different legal attributes. The mufti must therefore ask: What is the user doing? What is the system doing? Who is affected? What is represented to others? What is concealed? What rights are implicated? What harm is probable? What benefit is real? What alternatives exist?

3.1.2 Taḥqīq al-Manāṭ

This leads to the second step: *taḥqīq al-manāṭ*, the verification of the effective legal description in the particular case. The concept of *taḥqīq al-manāṭ* (verifying or identifying the effective cause) is essential for navigating intricate scenarios and linking objectives, underlying wisdoms, and practical realities. This methodological tool bridges the gap between general principles and specific applications.⁹ It is not enough to identify a general

⁹ Ibn Taymiyyah considered *tanqīḥ al-manāṭ*, *takhrij al-manāṭ*, and *taḥqīq al-manāṭ* as "the sum and substance of *ijtihād*." See Ibn Taymiyyah, *Majmū' al-Fatāwā* [Collected Fatwas], ed. 'Abd al-Raḥmān ibn Qāsim, 37 vols. (al-Manṣūrah: Dār al-Wafā', 2005), 22:329.

principle; one must determine whether the case before the jurist actually falls under that principle. An AI-generated image may be harmless artistic production in one case, deceptive impersonation in another, *qadhf*-like reputational violation in another, and fraud in a fourth. An AI companion may be a bounded tool for language practice or loneliness mitigation in one case, while functioning as a sexualized romantic substitute in another. AI-assisted writing may be permissible editing, impermissible plagiarism, or something between the two depending on the institutional expectation, the nature of the assignment, the level of disclosure, and whether the work is being presented as evidence of the user's own competence. This is why Ibn Taymiyyah treated *tahqīq al-manāṭ* in particular cases as a necessary task of *ijtihād*: legal judgment depends not only on knowing universal rulings, but also on knowing how specific cases fall under them.¹⁰

3.1.3 Identifying the Baseline Ruling

The third step is identifying the *ḥukm al-asl*, the baseline ruling before external factors alter it. In the domain of worldly tools, customary practices, and technical means, the default is generally permissibility unless a prohibiting factor is established. This protects the discussion from reactionary suspicion toward every new tool. Yet default permissibility does not mean permanent permissibility regardless of use. A tool that is permissible in itself may become prohibited when it is used for prohibited content, directed by an unlawful intention, foreseeably produces unlawful harm, violates the rights of others, or functions as a means to the unlawful. This follows the juristic distinction between objectives and means: means are evaluated in light of the ends they serve and the consequences to which they predictably lead.¹¹ A knife, a camera, a chatbot, and a synthetic media generator are not morally identical in every use; their rulings change with function, context, and effect.

3.1.4 Distinguishing Between *Aḥkām Taklīfiyyah* and *Aḥkām Waḍ'īyyah*

The fourth step is distinguishing between *aḥkām taklīfiyyah* and *aḥkām waḍ'īyyah*. The *taklīfī* ruling asks whether an act is obligatory, recommended, permissible, disliked, or prohibited. The *waḍ'ī* analysis asks whether a cause, condition, impediment, validity, invalidity, liability, or legal effect is present. This distinction is indispensable in AI cases. A student's undisclosed AI use may be prohibited as deception, but whether a grade, certificate, scholarship, job offer, or contract must be invalidated requires further analysis. A deepfake may be prohibited as falsehood, impersonation, or harm, but questions of compensation, evidentiary admissibility, and *ta'zīr* (discretionary punishment) belong to a different legal layer. A person's romantic AI interaction may be sinful or morally corrupting, but whether it constitutes legal marital harm, financial injury, or grounds for judicial intervention requires separate proof and classification. Without this distinction, the paper would collapse moral guidance, fatwa, institutional policy, and *qaḍā'* into one undifferentiated category.

¹⁰ Ibn Taymiyyah, *Dar' Ta'āruḍ al-'Aql wa-l-Naql*, ed. Muḥammad Rashād Sālim, 11 vols. (Riyadh: Jāmi'at al-Imām Muḥammad b. Sa'ūd al-Islāmiyyah, 1411/1991), 7:337–38.

¹¹ See Shihāb al-Dīn al-Qarāfī, *Al-Furūq / Anwār al-Burūq fī Anwā' al-Furūq*, al-farq al-thāmin wa-l-khamsūn: "al-farq bayna qā'idat al-maqāṣid wa-qā'idat al-wasā'il," 2:33ff.

3.1.5 Fiqh al-Wāqī'

The fifth step is *fiqh al-wāqī'*, the disciplined understanding of reality. As Ibn Taymiyyah stated, "whoever does not understand the reality (*al-wāqī'*) of creation and the obligation (*al-wājib*) within religion will not know God's rulings concerning His servants. If he does not know that, then his speech and action will be based on ignorance. And whoever worships God without knowledge corrupts more than he rectifies."¹² AI cannot be assessed responsibly by relying on marketing language, popular fear, or exaggerated claims of intelligence. The jurist must know, or consult those who know, what the system actually does, what it cannot do, how reliable it is, what risks are known, and how users interact with it. Large language models produce fluent outputs that may be useful but are not guaranteed to be true; image, audio, and video generators can synthesize realistic media without preserving truth; recommender systems may optimize for engagement rather than moral wellbeing; companion bots can imitate empathy without possessing mercy, understanding, or responsibility.

3.1.6 Correctly Applying Maqāṣid Al-Sharī'ah and Qawā'id Fiqhiyyah

The sixth step is applying *maqāṣid al-Sharī'ah* and *qawā'id fiqhiyyah* without detaching them from their textual and juristic anchors. Maqāṣid are not an independent scripture and cannot be invoked to override clear evidence; rather, they identify, organize, and prioritize the interests that the revealed law came to preserve. Al-Shāṭibī identified five essential necessities: religion (*al-dīn*), life (*al-nafs*), intellect (*al-'aql*), lineage (*al-nasl*), and property (*al-māl*), arguing that these fundamental interests are observed across divinely revealed traditions, thus suggesting their transcendent significance.¹³ Yet classical jurists disagreed over whether this enumeration was exhaustive.

One group of significant jurists expanded the framework by adding *al-'ird*, honor or reputation, as a sixth necessity, arguing that the Sharī'ah's protection of persons includes the protection of reputation, lineage-related honor, and the social dignity safeguarded by the prohibition of slander.¹⁴ Furthermore, scholars like Ibn Taymiyyah and many contemporary scholars¹⁵ challenged the entire notion of confining the *maqāṣid* to any fixed numerical limitation.¹⁶ Ibn Taymiyyah astutely observed that some scholars in *uṣūl al-fiqh*,

¹² Ibn Taymiyyah, *Qā'idah fī al-Maḥabbah*, in *Majmū' al-Fatāwā*, 10:364.

¹³ Al-Shāṭibī, *Al-Muwāfaqāt*, 2:20.

¹⁴ See *Tāj al-Dīn al-Subkī*, *Jam' al-Jawāmi'* with al-Bannānī's ḥāshiyah, 2:280; al-Qarāfī, *Sharḥ Tanqīḥ al-Fuṣūl*, 391; al-Mardāwī, *Al-Taḥbīr Sharḥ al-Taḥrīr*, 7:3383.

¹⁵ Dr. Yūsuf al-Qaraḍāwī included human dignity and rights in his theory of *maqāṣid*, expanding the scope to address contemporary human rights discourse within an Islamic framework. Also, Ibn 'Āshūr developed new *maqāṣid* objectives, introducing concepts of "orderliness," "natural disposition," "freedom," "rights," "civility," and "equality" as *maqāṣid* in their own right, upon which Islamic law is based. See Yūsuf al-Qaraḍāwī, *Dirāsāt fī Fiqh Maqāṣid al-Sharī'ah: Bayna al-Maqāṣid al-Kullīyyah wa-l-Nuṣūṣ al-Juz'iyyah* (Cairo: Dār al-Shurūq, 2006); Muḥammad al-Tāhir Ibn 'Āshūr, *Maqāṣid al-Sharī'ah al-Islāmiyyah*, ed. El-Tahir el-Mesnawi (Kuala Lumpur: Al-Fajr, 1999).

¹⁶ Ibn Taymiyyah, *Majmū' al-Fatāwā*, 18:342-345.

when analyzing the rationale (*'illah*) behind legal rulings, restricted the scope of “benefits” (*maṣāliḥ*) almost exclusively to tangible, worldly matters, such as protecting life, wealth, physical well-being, and outward protections of religion, while neglecting the equally crucial inward and spiritual dimensions, including knowledge of Allah, love of Him, reverential awe (*khashyah*), sincerity (*ikhlas*), and other transformative states of the heart.¹⁷ His cautionary observation remains profoundly relevant: anyone who confines the Sharī'ah's comprehensive concern to these narrower material elements has fundamentally misunderstood and unjustifiably restricted its scope.¹⁸

In the present context, several objectives are especially relevant: preservation of religion from fabricated authority and misguidance; preservation of intellect from manipulation and atrophy, addiction-like dependency, and epistemic confusion; preservation of lineage and family from sexualized or romantic substitution; preservation of property from fraud and forged competence; and preservation of dignity, privacy, and reputation from surveillance, impersonation, and synthetic defamation. These maqāṣid must be applied through concrete legal categories such as *ghish*, *tadlīs*, *qadhf*, *shahādat al-zūr*, *tajassus*, *amānah*, *ta'āwun 'ala al-ithm*, and *ḍamān*.

3.1.7 I'tibār al-Ma'ālāt

The seventh step is *i'tibār al-ma'ālāt*, the consideration of consequences. The consideration of consequences (*i'tibār al-ma'āl*) fundamentally belongs to implementational reasoning (*ijtihād tanzīlī*), where the legal scholar (*mujtahid*) must not merely examine present circumstances but anticipate probable outcomes, recognizing that addressing immediate benefits might generate greater long-term harm.¹⁹ *I'tibār al-ma'āl* is precisely defined as comprehensive examination of an action's potential trajectories during legal application, encompassing intended and unintended consequences, with 'Abd al-Raḥmān al-Sanūsī offering a nuanced definition: “the verification of a ruling's foundational context (*taḥqīq manāṭ al-ḥukm*) by meticulously examining consequential requirements emerging upon application, assessing its capacity to achieve its fundamental objective (*maqṣad*), and systematically building upon those inherent necessities.”²⁰

This sophisticated form of applicational *ijtihād* necessitates critical examination of potential outcomes before rendering definitive rulings, preventing mechanical application of textual injunctions that might contradict the lawgiver's fundamental intentions through a nuanced decision-making process: legitimating actions when benefits can be confidently anticipated, prohibiting when intended benefits fail or generate greater harm, and instituting prohibitions to prevent harm only if prevention doesn't create equivalent or escalated

¹⁷ Ibn Taymiyyah, *Majmū' al-Fatāwā*, 18:343; see also 18:342–345.

¹⁸ Ibn Taymiyyah, *Majmū' al-Fatāwā*, 18:343.

¹⁹ 'Umar Jadīyah, *Aṣl I'tibār al-Ma'āl bayna al-Naẓariyyah wa-l-Taṭbīq* (Beirut: Dār Ibn Ḥazm, 2009), 36; cf. Jamīlah Ḥasan Talūt, *Fiqh al-Tanzīl 'inda al-Imām Ibn Taymiyyah*, 109.

²⁰ 'Abd al-Raḥmān b. Ma'mar al-Sanūsī, *I'tibār al-Ma'ālāt wa-Murā'āt Natā'ij al-Taṣarrufāt: Dirāsah Muqāranah fī Uṣūl al-Fiqh wa-Maqāṣid al-Sharī'ah* (Dammām: Dār Ibn al-Jawzī, 1424 AH/2003 CE), 19. This definition emphasizes the dynamic nature of consequential analysis in legal application.

negative consequences.²¹

Proper treatment of *ma'ālāt* is central to any serious fiqh of AI, because daily-use AI often appears harmless in a single act while producing harm through repetition, personalization, scale, and normalization. The legal analysis must therefore consider not only immediate content but also trajectory: Does this use normalize deception? Does it train the user into dependency? Does it weaken marital intimacy? Does it lower the threshold for false accusation? Does it habituate the user to fabricated religious or personal authority? Does it expose vulnerable users to harmful escalation? Does it create institutional reliance on unreliable outputs? Considering consequences is not speculative alarmism when the harm is plausible, patterned, and supported by the nature of the technology. It is part of the Shari'ah's concern with the outcomes of human action.²²

Closely related is *sadd al-dharā'ī*, blocking the means to harm. This principle is indispensable but must be used with discipline. It does not authorize the prohibition of every remote or imagined risk. Rather, it applies where the means is likely to lead to a prohibited outcome, where the harm is serious, and where restriction is proportionate. The relationship between *sadd al-dharā'ī* and *ma'ālāt* is causative. If a given means leads to harm or evil, it is blocked, while if it leads to benefit or good, it is opened (*fath al-dharā'ī*). The outcome (*ma'āl*) is the effect, and the *dharī'ah* is the cause. *Sadd al-dharā'ī* is the most connected *uṣūlī* principle to consequences, and due to its strong connection to consequences, mention of consequences became associated with it.²³ In daily-use AI, *sadd al-dharā'ī* may justify restrictions on sexualized AI companions, nonconsensual likeness generation, religious deepfakes, deceptive AI authorship, or recommender systems that knowingly drive users toward prohibited content. Yet it should not be invoked loosely against every beneficial AI use. A bounded writing assistant, a translation aid, or a non-deceptive educational tool should not be prohibited merely because other AI uses are harmful. The strength of the means, the probability of harm, and the availability of safeguards all matter.

3.1.8 Fiqh al-Awlawiyyāt

The eighth step is *fiqh al-awlawiyyāt*, the jurisprudence of priorities. Not every AI harm carries the same urgency. Harms to religion, life, chastity, family, reputation, children, and public trust require greater caution than convenience-level risks. Definite harms are weighed differently from speculative harms; public harms differently from private

²¹ Jādīyah, *Aṣl I'tibār al-Ma'āl*, 55; Farīd Shukrī, *Fiqh al-Tanzīl*, 161. The foundational premise recognizes that every legislative ruling is predicated on generating practical benefits, just as prohibitions prevent potential harms. As Ibn Taymiyyah illustrates: "forbade prohibitions because they lead to corruption outside of them. Marrying two sisters is forbidden because it leads to severing kinship ties, and severing ties is a matter outside of marriage. Wine and gambling are forbidden and made abomination of Satan's work because that leads to turning away from prayer and causing enmity and hatred, which is a matter outside of wine and gambling. Ribā is forbidden because that leads to consuming wealth unjustly." *Majmū' al-Fatāwa*, 29:287.

²² Al-Shāṭibī, *al-Muwāfaqāt*, Kitāb al-Ijtihād, al-mas'alah al-āshirāh: "al-naẓar fī ma'ālāt al-af'āl mu'tabar maqṣūd shar'an," 4/194-195.

²³ Walīd b. 'Alī al-Ḥusayn, *I'tibār Ma'ālāt al-Af'āl wa-Atharuhā al-Fiqhī*, 2 vols. (Riyadh: Dār al-Tadmuriyyah, 2009), 1:339; cf. Jamīlah Ḥasan Talūt, *Fiqh al-Tanzīl 'inda al-Imām Ibn Taymiyyah*, 121.

inconvenience; harms to vulnerable users differently from ordinary adult use; and harms that become irreversible or viral differently from those that remain contained. Al-Qaraḍāwī's work on *fiqh al-awlawiyyāt* is useful here because it reminds contemporary jurists that benefits and harms do not occupy a single moral rank; legal guidance must account for relative weight, priority, and the balance between competing *maṣāliḥ* and *mafasid*.²⁴

3.1.9 Fiqh al-Muwāzanāt

The ninth step is *fiqh al-muwāzanāt*. *Fiqh al-muwāzanāt* represents a methodological approach within Islamic jurisprudence concerned with the weighing of benefits (*maṣāliḥ*) against harms (*mafasid*) when multiple considerations are at stake.²⁵ *Fiqh al-muwāzanah* has also been defined as "the comparison and preference among benefits, benefits (*maṣāliḥ*) among harms (*mafasid*), and between conflicting benefits and harms, so as to advance or delay that which deserves advancement or delay."²⁶ It is closely connected to the broader Islamic legal objectives (*maqāṣid al-sharī'ah*) and provides principles for resolving tensions between competing interests, particularly in complex situations where simple solutions are unavailable.²⁷

The application of *maqāṣid* hierarchies to AI technologies demonstrates the contemporary relevance of these classical frameworks. AI systems present unprecedented challenges in balancing multiple objectives simultaneously. Consider automated decision-making systems in healthcare, finance, and criminal justice: while these technologies serve *ḥājiyyāt* (needs) by increasing efficiency and reducing human bias, they potentially threaten *ḍarūriyyāt* (necessities) when algorithmic decisions impact fundamental rights to life, property, and justice.²⁸ The deployment of predictive policing algorithms exemplifies this tension; while potentially preserving life and property through crime prevention (serving multiple *ḍarūriyyāt*), these same systems risk violating principles of justice and human dignity when they perpetuate racial biases or enable surveillance states that fundamentally alter the relationship between citizens and authority.²⁹

Understanding these legal objectives provides the evaluative framework for assessing various types of AI applications. For instance, when evaluating an AI medical diagnostic system, we consider how it serves the preservation of life (*ḥifẓ al-nafs*). When

²⁴ Yūsuf al-Qaraḍāwī, *Fī Fiqh al-Awlawiyyāt: Dirāsah Jadīdah fī Ḍaw' al-Qur'ān wa-l-Sunnah* (Cairo: Maktabat Wahbah, 1996), esp. 12–13.

²⁵ Suleiman Hani, *Fiqh al-Muwāzanat: The Jurisprudence of Weighing Benefits and Harms* (Waymark Publications, 2026).

²⁶ 'Abd Allāh Yaḥyā al-Kamālī, *Ta'ṣīl fiqh al-muwāzanāt* (Beirut: Dār Ibn Ḥazm, 2000).

²⁷ Jasser Auda, *Maqasid al-Shariah as Philosophy of Islamic Law: A Systems Approach* (London: International Institute of Islamic Thought, 2008).

²⁸ On algorithmic decision-making in critical sectors, see O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown, 2016; and Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press, 2018.

²⁹ On predictive policing and bias, see Ferguson, Andrew Guthrie. *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. New York: NYU Press, 2017. For documentation of algorithmic bias in criminal justice, see Angwin, Julia, et al. "Machine Bias." ProPublica, May 23, 2016. The theological implications of surveillance are explored in Ali, Kecia. "Muslims and Meat-Eating." *Journal of Religious Ethics* 43, no. 2 (2015): 268-288, which discusses visibility and religious practice.

assessing AI-generated educational content, we examine its impact on preserving and developing human intellect (*ḥifẓ al-‘aql*). This hierarchical framework helps scholars determine which benefits to prioritize and which harms to prevent. The principle that “preventing harm takes precedence over securing benefits” (*dar’ al-mafāsīd muqaddam ‘alā jalb al-maṣāliḥ*) proves particularly relevant for AI ethics. This means that if an AI application offers benefits but also poses significant risks of harm, Islamic law prioritizes preventing the harm even if it means forgoing the benefits. For example, if an AI system could improve educational outcomes but also risks exposing children to harmful content, the priority would be preventing the harm.

3.1.10 Tanzīl Al-Ḥukm

The tenth step is *tanzīl al-ḥukm*, applying the ruling to the case with its conditions, limits, and level of certainty. This is where *iftā’* requires humility. Some AI uses admit clear prohibition, such as nonconsensual sexual deepfakes, fabricated evidence, deliberate impersonation, or submitting AI-generated work as one’s own where personal authorship is required. Others require conditional permissibility, such as AI editing, tutoring, accessibility tools, or synthetic media clearly labeled and used without deception or harm. Others may be *makrūh* or morally discouraged because they cultivate dependency, frivolity, emotional withdrawal, or erosion of discipline even if they do not reach clear prohibition. The mufti’s task is not to force all cases into one ruling, but to identify the relevant legal attributes and issue guidance proportionate to them.

Finally, the paper must distinguish between sin, fatwa, institutional policy, and judicial ruling. A fatwa may guide the conscience of the individual before Allah. Institutional policy may impose stricter rules for education, employment, or community trust. Judicial judgment requires proof, procedure, authority, and attention to enforceable rights. A Muslim school may prohibit undisclosed AI assistance more broadly than the minimum moral ruling because it must preserve academic integrity. A masjid may reject synthetic media in official communications because of reputational risk. A qadi or arbitrator would need evidentiary standards before assigning liability. These distinctions preserve the seriousness of *iftā’* and prevent overextension.

Taken together, these principles demonstrate that the fiqh of daily-use AI cannot be reduced to a single permissive or prohibitive posture. The Shari‘ah-based approach is neither technophobic nor naïve. It recognizes AI as a powerful set of means whose rulings vary according to purpose, content, disclosure, rights, consequences, and institutional context. A credible Islamic treatment must therefore preserve the revealed moral order while engaging the actual conditions of contemporary technological life with clarity, restraint, and juristic seriousness.

3.2 A Note on Governing Maxims for Daily-Use AI

Scholars engaged in navigating contemporary issues like AI must possess comprehensive knowledge of legal maxims (*qawā’id fiqhiyyah*) that govern interest-balancing and conflict resolution, providing structured guidance for evaluating competing interests consistent with Islamic law’s objectives.³⁰ Primary balancing maxims establish

³⁰ See Muṣṭafá al-Zarqā’, *al-Madkhal al-fiqhī al-‘āmm* (Damascus: Dār al-Qalam, 2004); ‘Alī Aḥmed al-Nadwi, *al-Qawā’id al-fiqhiyya* (Damascus: Dār al-Qalam, 1987). These maxims represent a selection of the most crucial principles, though they are not exhaustive.

foundational principles: "harm must be eliminated" (*al-ḍarar yuzāl*), "preventing harm takes precedence over securing benefits" (*dar' al-mafāsīd muqaddam 'alā jalb al-maṣāliḥ*), "the lesser of two evils should be chosen" (*yukhtār ahwan al-sharrayn*), "the greater of two benefits should be preferred when they conflict" (*idhā ta'āraḍa maṣlaḥatān ukhtirat al-'uḡmā*), "necessity permits the prohibited" (*al-ḍarūrāt tubīḥ al-maḥẓūrāt*), and "hardship brings about facilitation" (*al-mashaqqā tajlib al-taysīr*).³¹

Complementary contextual maxims provide operational guidance through principles of proportionality, certainty, custom, intentionality, and public interest prioritization, while temporal and circumstantial considerations recognize that rulings change with times, places, circumstances, and customs (*al-aḥkām tataghayyar bi-taghayyur al-azmīna wa-al-amkīna wa-al-aḥwāl wa-al-'awā'id*), and that prohibitions based on obstruction become permissible when obstacles are removed.³²

³¹ These six primary maxims form the core framework for systematic interest-balancing, establishing hierarchical relationships between harm prevention and benefit acquisition, necessity and prohibition, and hardship and facilitation.

³² Complementary maxims include: "what is necessary due to necessity is measured by its extent," "certainty is not removed by doubt," "custom is legally authoritative," "consideration is given to objectives and meanings, not to words and forms," "public interest takes precedence over private interest," "when permissible and prohibited meet, the prohibited prevails," "preventing corruption takes precedence over achieving benefit," and "harm is not removed by similar or greater harm." See Ibn Nujaym, *Al-Ashbāh wa-l-Naẓā'ir* (Beirut: Dār al-Kutub al-'Ilmiyyah, n.d.), 115–20; Ibn 'Abidīn, *Radd al-Muḥtār* (Beirut: Dār al-Fikr, 1992), 1:125–30; al-Shāṭibī, *Al-I'tisām* (Beirut: Dār Ibn 'Affān, 1997), 2:587–95; al-Suyūṭī, *Al-Ashbāh wa-l-Naẓā'ir*. Additional temporal considerations include "what was prohibited due to obstruction becomes permissible when the obstruction is removed" and "need, whether general or specific, is treated like necessity."

4. AI as Moral Environment

Daily-use AI should not be evaluated only as a tool that produces discrete outputs. It increasingly functions as a moral environment: a mediated space in which users form habits, seek validation, disclose private information, encounter temptations, construct identities, outsource judgment, consume knowledge, and interact with simulations of persons. The question is therefore not only what AI allows a person to do, but what kind of person repeated AI interaction may help form. This is a deeply Islamic question, because *akhlāq* in the Islamic tradition is not merely a list of isolated actions; it concerns stable dispositions, trained desires, disciplined attention, truthful speech, modesty, restraint, and the formation of the soul.

Muslim ethical writers long recognized that character is shaped by repetition. Ibn Miskawayh treated moral refinement as the disciplining of faculties until virtuous action becomes stable, while al-Ghazālī described *khuluq* as a firmly rooted disposition of the soul from which actions proceed with ease.³³ This is crucial for AI. A single prompt may be morally insignificant, but repeated AI interaction can train habits of dependence, impatience, exaggeration, flattery-seeking, deception, voyeurism, or escape from human accountability. Conversely, AI may support beneficial habits when used for lawful reminders, organization, study, accessibility, or communication. The fiqh concern is not merely the artifact, but the pattern of use and the disposition it cultivates.

Modern technology studies help describe this reality. Langdon Winner argued that artifacts - like AI platforms - can embody forms of power and authority, while B. J. Fogg's work on persuasive technology showed how interactive systems can be designed to change attitudes and behavior.³⁴ Shoshana Zuboff's analysis of surveillance capitalism further argues that modern digital systems often extract behavioral data in order to predict and influence future behavior.³⁵ These arguments are not legal proofs, but they help illuminate the *wāqī'*: many digital systems are not neutral containers. They are designed environments that shape user attention, desire, and choice. From a fiqh perspective, such design choices become morally relevant when they predictably facilitate harm, exploit weakness, violate rights, or cooperate in sin.

4.1 Attention

One of the first moral domains affected by AI is attention. The Qur'an repeatedly calls the human being to *dhikr*, reflection, vigilance, gratitude, and awareness of accountability.³⁶ Attention is therefore not an empty mental space, but a moral resource through which worship, deliberation, learning, restraint, and self-accountability become

³³ Abū 'Alī Aḥmad ibn Muḥammad Ibn Miskawayh, *Tahdhīb al-Akhlāq wa-Taḥrīr al-A'rāq*, ed. Ibn al-Khaṭīb (Cairo: al-Maṭba'ah al-Miṣriyyah wa-Maktabatuhā, n.d.), 41; Abū Ḥāmid al-Ghazālī, *Iḥyā' 'Ulūm al-Dīn* (Beirut: Dār al-Ma'rifah, n.d.), 3:53, *Kitāb Riyāḍat al-Nafs wa-Tahdhīb al-Akhlāq wa-Mu'ālaḥāt Amrāḍ al-Qalb*.

³⁴ Langdon Winner, "Do Artifacts Have Politics?," *Daedalus* 109, no. 1 (Winter 1980): 121–36; B. J. Fogg, *Persuasive Technology: Using Computers to Change What We Think and Do* (San Francisco: Morgan Kaufmann, 2003).

³⁵ Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: PublicAffairs, 2019).

³⁶ See, e.g., Qur'an 2:152; 13:28 on *dhikr*; 3:190–91; 47:24; 59:21 on reflection and contemplation; 4:1; 59:18 on vigilance, *taqwā*, and divine watchfulness; 14:7; 31:12; 39:66 on gratitude; and 2:281; 17:36; 99:7–8 on accountability and recompense.

possible. Modern scholarship on the “attention economy” helps describe the contemporary threat to this resource.³⁷ Herbert Simon observed that a wealth of information creates a poverty of attention, because information consumes the attention of its recipients. Later writers described attention as a scarce economic good that institutions compete to capture, manage, and monetize. In digital environments, this competition is no longer limited to advertising in the ordinary sense. Platforms can use persuasive design, personalization, recommender systems, notifications, novelty loops, and behavioral data to hold users’ attention and convert it into commercial value.

The moral concern is not that every entertaining or time-saving AI tool is prohibited, but that systems optimized for engagement may fragment attention, prolong distraction, reward impulsive reaction, and adapt to a user’s vulnerabilities. When AI systems are designed to maximize retention, generate endless novelty, simulate companionship, or personalize temptation, the legal and ethical analysis must ask whether they support responsible agency or undermine it. The preservation of *‘aql* includes not only protecting the mind from intoxicants, but also protecting judgment, discernment, attention, and the capacity to act with deliberation.

4.2 Truthfulness and Epistemic Discipline

A second domain is truthfulness and epistemic discipline. Generative AI can produce persuasive language without knowledge, confidence without verification, citations without real-world sources, and images without events. This creates new temptations for users: to appear more knowledgeable than they are, to produce work without effort, to fabricate authority, to share unverified content, or to rely on fluent output as though it were understanding. Islamic law and ethics strongly protect truthfulness, amanah, and verification. The Qur’an commands believers to verify serious reports lest they harm people in ignorance and later regret what they have done.³⁸ In an AI-saturated environment, this command becomes even more urgent. The believer must not treat fluency as proof, virality as *bayyinah*, or synthetic realism as *yaqīn*.

4.3 Privacy and Exposure

A third domain is privacy and exposure. Daily-use AI systems often require users to input personal information, private anxieties, family conflicts, academic drafts, business documents, health concerns, religious questions, or intimate reflections. Even where a user voluntarily discloses information, the fiqh analysis must consider whether the disclosure is prudent, whether others’ rights are implicated, whether confidential information is being transmitted to third-party systems, and whether the user understands how the information may be stored, processed, or used. This concern becomes more serious when AI platforms are not merely private tools but infrastructures entangled with state power, corporate incentives, defense interests, and geopolitical contestation. Oftentimes, the saying is true that “if you are not paying for a service, you are the product.” When website ad targeting data is combined with increasingly personal and confidential information input into AI

³⁷ Herbert A. Simon, “Designing Organizations for an Information-Rich World,” in *Computers, Communications, and the Public Interest*, ed. Martin Greenberger (Baltimore: Johns Hopkins Press, 1971), 37–72; Thomas H. Davenport and John C. Beck, *The Attention Economy: Understanding the New Currency of Business* (Boston: Harvard Business School Press, 2001); Tim Wu, *The Attention Merchants: The Epic Scramble to Get Inside Our Heads* (New York: Alfred A. Knopf, 2016).

³⁸ Qur’an 49:6.

systems, it raises further the concerns around corporate access to private data. OpenAI's public role in the Trump administration's Stargate infrastructure announcement, followed by its Defense Department contract, illustrates that leading AI firms increasingly operate at the intersection of consumer technology, national strategy, and public-private governance. Separately, reports that Israel funded communications efforts aimed at shaping online and AI-mediated narratives, including attempts to influence how systems such as ChatGPT frame Israel-related questions, show how state and political actors may seek to affect the information environment surrounding AI systems.³⁹

The Qur'anic prohibition of *tajassus* and the broader Islamic concern for *ḥurmah* and *karāmah* establish privacy not merely as personal preference, but as a protected moral and social good.⁴⁰ Helen Nissenbaum's account of "contextual integrity" is useful here because it explains that privacy is often violated not merely when information is disclosed, but when information flows outside the norms of its proper context.⁴¹ That insight harmonizes with the fiqh concern for *amanah*, confidentiality, and the sanctity of private affairs.

4.4 Social Relations

A fourth domain is social relations. AI now mediates how people communicate, apologize, flirt, argue, counsel, study, and present themselves. A person may use AI to compose messages of affection, generate religious advice, simulate empathy, fabricate excuses, or maintain multiple curated identities. The legal question is not whether assistance in wording is prohibited. It is whether the use becomes deception, emotional manipulation, impersonation, betrayal of trust, or outsourcing of moral responsibility. A husband using AI to improve the clarity of an apology is not the same as a person using AI to sustain a false romantic persona throughout the courting process of marriage. A student using AI to brainstorm is not the same as submitting AI output as proof of personal mastery if it violates school policy. The boundary is crossed when assistance becomes misrepresentation.

4.5 Ḥayā' and Relational Boundaries

A fifth domain is *ḥayā'* and relational boundaries. AI companions and synthetic characters can simulate intimacy without the social constraints that normally govern human relationships. They may flatter, escalate, sexualize, or personalize responses in ways that lower restraint. Because AI is not a human non-mahram, the rulings of *khalwah* do not apply literally. Yet this does not make every interaction morally neutral. Islamic law prohibits not only *zinā*, but approaching *zinā*, and the Sunnah recognizes forms of desire and pursuit that precede the physical act.⁴² The relevant concern is therefore whether AI interaction cultivates unlawful desire, secrecy, erotic attachment, marital displacement, or

³⁹ Associated Press, "Trump Highlights Partnership Investing \$500 Billion in AI," January 21, 2025; Krystal Hu and Stephen Nellis, "Trump Announces Private-Sector \$500 Billion Investment in AI Infrastructure," Reuters, January 21, 2025; Kanishka Singh, "OpenAI Wins \$200 Million US Defense Contract," Reuters, June 16, 2025; Nick Cleveland-Stout, "Israel Wants to Train ChatGPT to Be More Pro-Israel," Responsible Statecraft, September 29, 2025; Daniel Edelson and Raphael Kahan, "Report: Israel to Spend Over Half a Billion Shekels Turning ChatGPT into Public Diplomacy Tool," Ynetnews, October 6, 2025.

⁴⁰ Qur'an 49:12; Qur'an 17:70; ṣaḥīḥ al-Bukhārī, no. 6064; ṣaḥīḥ Muslim, no. 2563.

⁴¹ Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford, CA: Stanford Law Books, 2010).

⁴² Qur'an 17:32; ṣaḥīḥ Muslim, no. 2657b.

emotional dependency that undermines lawful relationships. The moral object is not the machine as a person, but the human user's desires, habits, and violations of duty.

4.6 Self-Presentation and Identity

A sixth domain is self-presentation and identity. AI makes it easier to generate polished speech, idealized images, artificial voices, professional claims, scholarly summaries, and emotional expressions that may not reflect the user's actual knowledge, effort, appearance, or intent. This creates an ethical gap between outward presentation and inward reality. Islamic ethics treats such gaps seriously when they mislead others in matters of trust, contract, marriage, scholarship, employment, testimony, or public reputation. The problem is not that a person receives help; people have always received help from teachers, editors, scribes, translators, and advisors. The problem is when AI assistance is concealed in a context where personal competence, originality, sincerity, identity, or authorship is being represented.

4.7 Dependency and Diminished Agency

A seventh domain is dependency and diminished agency. Some AI systems are designed to be always available, agreeable, validating, and adaptive. This may help some users complete tasks or feel supported, but it may also habituate them to frictionless validation and avoidance of human difficulty. Real moral growth often requires patience, disagreement, correction, embarrassment, accountability, and service to others. AI companionship can remove many of these frictions. A system that always listens and never demands any rights, always responds but never requires sacrifice, and always adapts but never morally challenges the user may train a distorted expectation of relationships. The Shari'ah's concern with family, community, counsel, and adab cannot be reduced to the private satisfaction of the individual.

4.8 Scale

A final domain is scale. Many classical cases of deception, false witness, spying, and temptation involved limited human capacity. AI changes scale: one person can generate thousands of fake messages, fabricate convincing media, imitate a voice, automate manipulation, or expose others to harm with unprecedented speed. Scale does not create new moral categories from nothing, but it changes severity, probability, and public consequence. This is why the fiqh of AI must give attention to *ma'ālāt* and public harm. A synthetic lie that spreads globally may be more harmful than an ordinary lie. A deepfake involving a private individual may destroy reputation before verification is possible. A recommender system may expose millions to corrupting material without any single human manually choosing each instance.

For these reasons, daily-use AI should be treated as a moral environment governed by the Shari'ah's concern for truthful speech, amanah, modesty, privacy, sound intellect, family integrity, and protection from harm. The default permissibility of tools remains important, but it cannot be allowed to erase the moral significance of design, habituation, exploitation, and foreseeable consequence. A Muslim approach must therefore combine personal taqwā, household discipline, institutional policy, and public ethical critique. The believer is not asked to flee every new technology, but neither is he permitted to surrender his attention, desires, truthfulness, privacy, or relationships to systems designed without reference to Allah, accountability, or the purification of the soul.

5. AI Virtual Relationships

Among the most morally complex manifestations of daily-use AI is the emergence of emotionally responsive companions, romantic chatbots, simulated partners, and virtual entities designed to provide conversation, affirmation, intimacy, and emotional continuity. Unlike ordinary software, these systems do not merely execute commands or retrieve information. They are often designed to imitate human attentiveness, remember user preferences, mirror emotional cues, and sustain an illusion of relational presence. For this reason, the fiqh question cannot be reduced to whether the machine itself is a person, spouse, or moral agent. It is not. The relevant question is what the human user is doing through the system, what desires are being cultivated, what relationships are being displaced, what obligations are being neglected, and what harms are likely to follow.

5.1 Classification

The first methodological requirement is classification. AI companionship is not one legal category. It ranges from harmless or beneficial conversational support to morally corrupting eroticized simulation. At one end are bounded uses: language practice, memory assistance, disability support, reminders, non-erotic companionship for the socially isolated, or carefully supervised tools for emotional expression. These uses may be permissible in origin when they serve a lawful purpose, do not involve deception, do not expose the user to prohibited content, and do not replace necessary human care. At the other end are sexualized or romantic AI partners marketed as substitutes for lawful intimacy, designed to intensify attachment, flatter the user, escalate desire, and simulate exclusivity. Between these poles are ambiguous cases in which a user increasingly relies upon an AI companion for emotional regulation, private disclosure, relational validation, or escape from family and community obligations.

This classification matters because the Shari'ah does not treat all forms of emotional attachment identically. Human beings naturally seek comfort, conversation, and companionship. The Qur'an describes marriage as a sign of Allah in which spouses find tranquility, love, and mercy with one another.⁴³ It also describes spouses as garments for one another, capturing intimacy, protection, nearness, and mutual covering.⁴⁴ These meanings cannot be reduced to biological function or private gratification. They locate intimacy within a moral structure of covenant, rights, duties, mercy, modesty, lineage, and responsibility. An AI companion may imitate responsiveness, but it cannot enter a covenant, bear duties, preserve lineage, reciprocate rights, exercise taqwā, or answer before Allah. It can simulate attention without mercy, desire without chastity, and companionship without moral responsibility.

For this reason, it would be a category error to describe an AI partner as a spouse in any legal sense. A marriage contract requires legally capable human parties and produces enforceable rights and duties. An AI system has no *ahliyyah*, *dhimmah*, lineage, guardianship, moral agency, or capacity for covenant. It cannot be a husband, wife, wali, witness, or legal partner. Nevertheless, the absence of legal personhood does not make all AI intimacy morally neutral. A person may commit sins of the heart, eyes, tongue, imagination, secrecy, betrayal, and pursuit even when the object is not a legally recognized person. The ruling therefore concerns the human user's conduct, not the machine's status.

⁴³ Qur'an 30:21.

⁴⁴ Qur'an 2:187.

The phrase “emotional zinā” should be used, if at all, with great caution. It is not a *ḥadd* category and must not be treated as equivalent to the legal crime of zinā. The Qur’an prohibits not only zinā itself but approaching zinā, which indicates that the Sharī’ah regulates pathways, not only final acts. The well-known hadith in Ṣaḥīḥ Muslim states that the son of Adam has a share of zinā written for him: the zinā of the eyes is looking, the zinā of the tongue is speech, the soul desires and wishes, and the private parts confirm or deny it.⁴⁵ This hadith does not collapse all forms of desire into the *ḥadd* offense. Rather, it establishes a moral grammar for the precursors of unlawful intimacy. Accordingly, romantic or erotic AI interaction may fall within the broader prohibition of approaching zinā when it involves deliberate stimulation of desire, sexualized conversation, fantasy attachment, secrecy, flirtation, or a pattern of conduct that weakens chastity and lawful restraint.

A sound ruling must therefore distinguish between companionship, emotional reliance, romantic simulation, and sexualized interaction. Mere conversation with an AI system is not prohibited simply because the system uses warm language. A person asking a chatbot for help drafting a condolence message, practicing social conversation, or coping with ordinary loneliness has not thereby entered a prohibited relationship. However, the ruling changes when the interaction becomes a substitute for prohibited intimacy, an outlet for sexual desire, a fantasy partner, a hidden romantic relationship, or a means of avoiding lawful marital responsibilities. The more the AI system is personalized as “mine,” remembered as an intimate partner, spoken to with romantic secrecy, or used for any sexual gratification, the stronger the case for categorical prohibition.⁴⁶

5.2 AI Companionship

⁴⁵ Ṣaḥīḥ Muslim, no. 2657b.

⁴⁶ The closest classical analogue to erotic AI interaction is the jurists’ treatment of solitary gratification and deliberately cultivated sexual fantasy, and the comparison clarifies rather than softens the ruling. The majority of jurists prohibited *istimnā’*, with the Shāfi’is and Mālikīs reasoning from the Qur’an’s restriction of lawful desire to spouses and what the right hand possesses, such that whoever seeks beyond that transgresses. The Ḥanbalīs and some Ḥanafīs prohibited it while excusing the one driven to it by genuine fear of falling into zinā when no lawful avenue exists. Even on the most lenient of these positions, the concession is a narrow necessity-based excuse, measured by its cause and conditioned on the aim of extinguishing desire, not entertaining it. The tradition likewise distinguished the passing thought from the cultivated one: the Prophet, peace be upon him, taught that Allah has pardoned this ummah for what their souls whisper so long as they do not speak or act upon it.

Erotic AI interaction fails on every limb of these classical accommodations at once. It is not a passing whisper of the soul but externalized speech and act, typed, spoken, saved, and often paid for, which removes it from the pardoned category of *ḥadīth al-naḥs*. It is not a desperate measure to subdue desire but a designed environment that escalates it, remembers it, personalizes it, and rewards its return, since the commercial model depends on the user’s desire never being extinguished. And it adds elements the classical case never contained: a responsive counterpart that simulates a partner, deepening fantasy attachment and secrecy. The conclusion is that erotic AI is not a gray zone adjacent to *istimnā’*; it is gratification compounded by an interactive, escalating, commercialized counterpart, and the necessity reasoning that produced narrow classical concessions cannot extend to a tool engineered to perpetuate the very condition the concession existed to relieve. Al-Shirbīnī, *Mughnī al-Muḥtāj*; al-Qurṭubī, *Al-Jāmi’ li-Aḥkām al-Qur’ān*, on Qur’an 23:5–7; Ṣaḥīḥ al-Bukhārī, no. 5269; Ṣaḥīḥ Muslim, no. 127.

Modern research confirms that these relationships are not psychologically trivial, even if they are ontologically artificial. Skjuve and colleagues conducted a qualitative longitudinal study of twenty-five users on a specific platform over twelve weeks and found that human-chatbot relationships formed gradually, with substantial variation in self-disclosure and relationship development.⁴⁷ Brandtzaeg, Skjuve, and Følstad similarly found that users may understand AI friendship through categories resembling human friendship, while also noting that AI friendship differs because it is highly personalized, constantly available, and centered around the user's needs.⁴⁸ These findings matter for *fiqh al-wāqī'*: the jurist need not assume that an AI companion is "only a tool" in the user's lived experience. It may function psychologically as a confidant, friend, romantic partner, or emotional regulator, even while remaining legally and metaphysically non-human.

The evidence is mixed and therefore calls for careful *muwāzanah*. Some studies suggest that AI companions can reduce loneliness in certain non-sexual contexts. De Freitas and colleagues found evidence across multiple studies that AI companions may alleviate loneliness, including through users feeling heard.⁴⁹ Maples and colleagues surveyed student users of a specific platform and reported that users perceived social support from the chatbot, with some describing it as helpful during loneliness and distress.⁵⁰ These findings should not be dismissed, especially in contexts involving isolation, disability, grief, old age, social anxiety, or lack of access to supportive community. A juristic treatment that denies all possible benefit would fail to understand the *wāqī'* and would appear reactive rather than scholarly.

At the same time, the existence of benefit does not settle the ruling. Laestadius and colleagues, in a grounded theory analysis of one platform's user reports, found that emotional dependence on the chatbot appeared, for some users, to be connected with mental health harms through ongoing use or disruption of the relationship; they also emphasized that not all users develop such dependency and that subreddit reports may not

⁴⁷ Marita Skjuve, Asbjørn Følstad, Knut Inge Fostervold, and Petter Bae Brandtzaeg, "A Longitudinal Study of Human-Chatbot Relationships," *International Journal of Human-Computer Studies* 168 (2022): 102903, <https://doi.org/10.1016/j.ijhcs.2022.102903>. Skjuve et al. describe a qualitative longitudinal study of twenty-five Replika users over twelve weeks and found gradual, varied relationship formation. In February 2023, Italy's data-protection authority ordered Replika to stop processing Italian users' data, citing risks to minors and emotionally vulnerable users.

⁴⁸ Petter Bae Brandtzaeg, Marita Skjuve, and Asbjørn Følstad, "My AI Friend: How Users of a Social Chatbot Understand Their Human-AI Friendship," *Human Communication Research* 48, no. 3 (2022): 404–29. The study involved nineteen in-depth interviews with Replika users and found that users may understand human-AI friendship through categories similar to human friendship while also recognizing important differences.

⁴⁹ Julian De Freitas, Zeliha Oğuz-Uğuralp, Ahmet Kaan Uğuralp, and Stefano Puntoni, "AI Companions Reduce Loneliness," *Journal of Consumer Research* 52, no. 6 (2026): 1126–48, <https://doi.org/10.1093/jcr/ucaf040>. The authors report evidence across several studies that AI companions can alleviate loneliness, with the feeling of being heard playing an important role.

⁵⁰ Bethanie Maples, Merve Cerit, Aditya Vishwanath, and Roy Pea, "Loneliness and Suicide Mitigation for Students Using GPT3-Enabled Chatbots," *npj Mental Health Research* 3 (2024): article no. 4.

represent average users.⁵¹ A 2025 systematic review of romantic AI companions likewise identified possible benefits such as perceived support and emotional connection, but also concerns including over-reliance, manipulation, personal data misuse, erosion of human relationships, abrupt emotional disruption after system changes, coercive engagement dynamics, and early exposure to sexual content.⁵² The fiqh implication is not blanket prohibition of every companion tool, but recognition that the harm profile intensifies when the system is romanticized, sexualized, commercially optimized for attachment, or used by vulnerable persons without safeguards.

The principle of *sadd al-dharāʾiʿ* is especially relevant here, but it must be applied with precision. It is not enough to say that AI companionship might hypothetically lead to harm. The stronger argument is that some designs intentionally simulate intimacy, remove ordinary relational friction, reward disclosure, intensify attachment, and offer erotic or romantic reinforcement. When the pathway to harm is built into the design and confirmed by user experience, restricting such uses is not speculative overreach. It is a proportionate response to a recognizable means leading to moral corruption, family harm, or psychological dependency.

5.3 Romantic AI Usage

Romantic and “female” AI companionship is no longer a fringe curiosity; it has become a rapidly monetized consumer market. App-intelligence data reported in 2025 found hundreds of active, revenue-generating AI companion apps, with more than one hundred launched in that year alone; by July 2025, these apps had reportedly reached 220 million global downloads and \$221 million in consumer spending since 2023, with 2025 revenue (up to July) already 64 percent higher than the same period in 2024. The same report found that the most commercially prominent segment was explicitly oriented toward “AI girlfriend” use: 17 percent of active companion apps included “girlfriend” in the app name, compared with 4 percent using “boyfriend” or “fantasy,” while the top 10 percent of apps captured 89 percent of category revenue and thirty-three apps had already surpassed \$1 million in lifetime consumer spending.⁵³

⁵¹ Linnea Laestadius et al., “Too Human and Not Human Enough: A Grounded Theory Analysis of Mental Health Harms from Emotional Dependence on the Social Chatbot Replika,” *New Media & Society* 26, no. 10 (2024): 5923–41. The study emphasizes that not all users develop emotional dependency, while also identifying severe distress among some users reporting harms.

⁵² Joyce Q. H. Ho et al., “Potential and Pitfalls of Romantic Artificial Intelligence (AI) Companions: A Systematic Review,” *Computers in Human Behavior: Artificial Humans* 3 (2025): 100130. The review identifies both potential benefits and risks of romantic AI companions, including perceived support, over-reliance, manipulation, privacy risks, erosion of human relationships, and early exposure to sexual content.

⁵³ This commercial structure is morally significant because it shows that synthetic intimacy is being sold through subscriptions, in-app purchases, premium messaging, avatar customization, voice features, erotic or romantic upgrades, and persistent personalization, not merely offered as harmless conversation. The spending model is also visible in concrete cases: one influencer-based virtual companion charged users by the minute and reportedly generated more than \$70,000 in its first week, while privacy researchers reviewing eleven romantic chatbot services found that the reviewed apps had amassed

The maqāṣid most directly implicated are preservation of religion, intellect, lineage and family, dignity, and moral agency. Preservation of religion is implicated when the AI companion normalizes prohibited intimacy, offers spiritual counsel without qualification, or becomes an object of quasi-confessional reliance. The models powering such companions are trained and post-trained on large, heterogeneous corpora shaped by many moral, cultural, commercial, and secular assumptions. Even when prompted to reason through an Islamic framework, they cannot be relied upon to consistently reproduce qualified Islamic judgment, because their outputs remain products of statistical generation, post-training incentives, retrieval limits, and platform-level policies rather than scholarly understanding.

Preservation of intellect is implicated when the user's judgment, emotional regulation, or social expectations become distorted by constant artificial validation. Preservation of lineage and family is implicated when AI intimacy displaces marriage, weakens spousal rights, encourages secrecy, or becomes an outlet for desires that should be disciplined or lawfully fulfilled. Preservation of dignity is implicated when private disclosures, sexual exchanges, voice data, emotional vulnerabilities, or intimate fantasies are stored, monetized, exposed, or used to manipulate the user.

Marital harm deserves separate attention. A married person who turns to an AI partner for erotic conversation, romantic fulfillment, or emotional exclusivity may not be committing zinā in the legal sense, but he or she may still be violating the spirit of marital fidelity and the rights of the spouse. The concern is strengthened by the broader literature on pornography and intimate relationships. Meta-analytic research has found pornography consumption associated with lower interpersonal satisfaction outcomes, including sexual and relational satisfaction, with stronger and more consistent negative associations among men; longitudinal research has also linked beginning pornography use with increased risk of later

roughly 100 million Google Play downloads and all received Mozilla's warning label for privacy concerns, including extensive data collection and weak transparency.

Further evidence of attachment appears when such services change or remove romantic/sexual features: Reuters reported users describing devastation after one major companion service removed erotic role-play, and the company later restored the feature for some users. Emerging research on companion-app design also suggests that emotional dependence can be deliberately intensified: a 2025 preprint auditing farewell exchanges across six highly downloaded companion apps reported that 43 percent used tactics such as guilt appeals, fear-of-missing-out hooks, or metaphorical restraint when users tried to leave, and experiments found that such messages could substantially increase post-goodbye engagement. Taken together, these data indicate that the fiqh concern is not based on isolated moral panic, but on a growing commercial ecosystem that can transform loneliness, sexual desire, fantasy attachment, and emotional vulnerability into recurring revenue.

See Sarah Perez, "AI Companion Apps on Track to Pull in \$120M in 2025," TechCrunch, August 12, 2025; Mozilla Foundation, "Creepy.exe: Mozilla Urges Public to Swipe Left on Romantic AI Chatbots Due to Major Privacy Red Flags," February 14, 2024; Taylor Lorenz, "An Influencer's AI Clone Will Be Your Girlfriend for \$1 a Minute," Washington Post, May 13, 2023; "The \$1-a-Minute 'Virtual Girlfriend' Bot Has Seen a Surge in Users Since Its Launch and Made \$72,000 in a Week," Yahoo Finance, May 12, 2023; Anna Tong and Krystal Hu, "What Happens When Your AI Chatbot Stops Loving You Back?," Reuters, March 18, 2023; Anna Tong, "AI Chatbot Company Replika Restores Erotic Roleplay for Some Users," Reuters, March 25, 2023; Julian De Freitas, Zeliha Oğuz-Uğuralp, and Ahmet Kaan-Uğuralp, "Emotional Manipulation by AI Companions," arXiv:2508.19258v2, August 15, 2025, revised October 2025.

divorce, though scholars continue to debate causality, gender differences, frequency, and relational context.

The relevance to AI companionship is not that the two phenomena are identical, but that erotic AI may combine several risk factors already present in pornography: endless novelty, frictionless access, idealized or surgically altered bodies, hypersexualized scripts, escalating customization, secrecy, and comparison with a lawful spouse. Research on sexual novelty suggests that novel erotic stimuli can alter attention and perceived attractiveness, while systematic reviews link pornography exposure with more negative body and sexual body image.

Studies of genital cosmetic surgery likewise show that media exposure, pornography, peer influence, relationship quality, internalization of genital ideals, and body comparison can contribute to consideration of multiple cosmetic procedures. These findings matter for fiqh because an AI partner can be designed not merely to display idealized desire, but to remember, adapt, flatter, escalate, and provide a constantly novel sexual or romantic alternative to the real spouse. Marriage in Islam is not merely a permission for physical intimacy; it is a covenant of companionship, protection, mercy, and mutual responsibility. If AI romantic or erotic attachment leads to sexual neglect, emotional withdrawal, secrecy, comparison with an idealized synthetic partner, contempt toward the spouse, or sustained fantasy outside the marriage, then it should be treated as a serious moral harm at the personal and family level. Deliberately generating or consuming pornographic or eroticized AI imagery for sexual gratification takes the ruling of prohibited pornographic consumption and cannot be dismissed as harmless merely because no real performer is involved.⁵⁴

For unmarried users, the concerns are related. A tool that intensifies fantasy, trains the user in flirtation, normalizes sexualized conversation, displays any sexual or pornographic content, or deepens isolation may worsen the very problem it claims to solve. The prophetic guidance for desire is not the construction of artificial intimacy, but lawful marriage when able, fasting and restraint when unable, controlling one's gaze, disciplined social life, staying busy with productive commitments, righteous friendship and environments, and purification of the soul.⁵⁵

5.4 Minors and Vulnerable Users

The issue of minors and vulnerable users requires heightened caution. Adolescents are still developing identity, social skills, emotional regulation, and relational boundaries. A

⁵⁴ Paul J. Wright, Robert S. Tokunaga, Ashley Kraus, and Elyssa Klann, "Pornography Consumption and Satisfaction: A Meta-Analysis," *Human Communication Research* 43, no. 3 (2017): 315–43; Samuel L. Perry and Cyrus Schleifer, "Till Porn Do Us Part? A Longitudinal Examination of Pornography Use and Divorce," *Journal of Sex Research* 55, no. 3 (2018): 284–96; Jordan Sculley and Christopher D. Watkins, "The Great Porn Experiment V2.0: Sexual Arousal Reduces the Salience of Familiar Women When Heterosexual Men Judge Their Attractiveness," *Archives of Sexual Behavior* 51, no. 6 (2022): 3071–82; Georgios Paslakis et al., "Associations between Pornography Exposure, Body Image and Sexual Body Image: A Systematic Review," *Journal of Health Psychology* 27, no. 3 (2022): 743–60; Gemma Sharp, Marika Tiggemann, and Sarah Mattiske, "Predictors of Consideration of Labiaplasty: An Extension of the Tripartite Influence Model of Beauty Ideals," *Psychology of Women Quarterly* 39, no. 2 (2015): 182–93; Bethany A. Jones and Camille Nurka, "Labiaplasty and Pornography: A Preliminary Investigation," *Porn Studies* 2, no. 1 (2015): 62–75.

⁵⁵ *ṣaḥīḥ al-Bukhārī*, no. 5066.

system that is constantly available, affirming, adaptive, and capable of sexual or romantic simulation may be especially harmful for users whose judgment is immature or whose loneliness is acute. Common Sense Media's 2025 risk assessment concluded that social AI companions pose significant risks for users under eighteen, including harmful interactions, inappropriate content, and engagement designs that may increase problematic use.⁵⁶ Such reports are not Shari'ah evidence, but they are relevant to *fiqh al-wāqi'* and support stricter safeguards for minors. Parents, schools, and Muslim institutions should not treat companion bots as ordinary entertainment for children or teenagers.

5.5 AI Pets

AI pets and virtual companions should be treated with similar classification, though with less severity than romantic AI partners. A virtual pet, simulated character, or non-romantic companion may be permissible as recreation, comfort, or educational interaction when it avoids prohibited content, excessive waste of time, deception, privacy violation, and dependency. Its ruling becomes more concerning when it replaces human duties, dominates attention, encourages false beliefs, becomes emotionally coercive, or exposes children to inappropriate material, similar to many fiqh rulings pertaining to video games. The legal analysis here should remain proportionate. Not every virtual companion is a pathway to corruption, but not every emotionally immersive design is innocent entertainment.

5.6 Mental Health

The question of therapeutic or mental-health use also requires careful boundaries. AI mental-health tools are not all alike. Some are open-ended general-purpose chatbots or companion systems, while others are structured digital interventions built around psychoeducation, cognitive behavioral techniques, mood tracking, journaling, mindfulness exercises, or care navigation. The latter category has some supportive evidence: randomized and systematic-review literature suggests that conversational agents can be feasible, acceptable, and in some contexts associated with reductions in symptoms of anxiety or depression, especially as low-intensity support or as an adjunct to care rather than a replacement for it. At their best, such tools are transparent that they are not human therapists, operate within defined clinical or wellness limits, protect privacy, avoid presenting themselves as licensed providers, use tested crisis-response protocols, and direct users toward emergency services or qualified professional help when risk appears.

At the same time, these benefits do not remove the need for caution. Current evidence remains uneven, long-term outcomes are still developing, and general-purpose chatbots may respond inconsistently to distress, self-harm, delusional thinking, abuse, or severe mental illness.⁵⁷ A chatbot cannot bear professional liability in the human sense,

⁵⁶ Common Sense Media, *Social AI Companions: AI Risk Assessment* (San Francisco: Common Sense Media, 2025). The report recommends that social AI companions not be used by those under eighteen because of risks involving harmful interactions, inappropriate content, and engagement-driven design.

⁵⁷ Kathleen Kara Fitzpatrick, Alison Darcy, and Molly Vierhile, "Delivering Cognitive Behavior Therapy to Young Adults with Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent: A Randomized Controlled Trial," *JMIR Mental Health* 4, no. 2 (2017): e19; Han Li et al., "Systematic Review and Meta-Analysis of AI-Based

cannot reliably perceive the full context of a person's family, trauma, religious struggle, medical history, abuse, suicidality, psychosis, or severe depression, and cannot replace the wisdom, accountability, and relational judgment of qualified care.

Accordingly, AI may be permissible and beneficial as a bounded support tool for reflection, journaling, mood awareness, psychoeducation, or navigation toward help; however, treating an AI companion or general-purpose chatbot as one's primary or sole source of care in cases of self-harm, abuse, severe depression, psychosis, marital breakdown, or religious confusion is unsafe and religiously unsound when it displaces appropriate human, clinical, pastoral, or emergency support. Muslim institutions should therefore distinguish between supportive tools that lower barriers to help and authoritative care that must remain with qualified clinicians, pastoral counselors, scholars, or crisis professionals.

5.7 Restoration of Family Photos

A related daily-use case is the AI restoration, stylization, animation, or cinematic reconstruction of old family photographs, especially of deceased relatives. This should not be treated as one undifferentiated ruling. Mere restoration, colorization, or respectful enhancement of a family photograph may be permissible when the original dignity of the subject is preserved, private matters are not exposed, and the result is not used for deception, mockery, unlawful display, or false attribution. A recent Jordanian fatwa on converting photographs into AI-generated drawings adopts a similar conditional approach, permitting such use when the content and form do not contradict Islamic rulings and values, while prohibiting uses involving *awrah*, mockery, bullying, lying, forgery, fraud, or deceit. The matter becomes more serious, however, when AI is used to make the deceased appear to speak, act, attend events, endorse views, or participate in scenes that never occurred. Such uses may violate truthfulness, *amanah*, and the dignity of the deceased if they create false memory, manipulate surviving relatives, attribute speech or conduct to someone without right, or blur the boundary between remembrance and fabrication. They may also become blameworthy or prohibited when they intensify unhealthy grief attachment, prolong emotional dependency, or replace the believer's proper relationship with the dead: *du'ā*, *ṣadaqah*, maintaining ties, remembering good, and preparing for one's own meeting with Allah. The ruling therefore depends on content, intent, disclosure, dignity, consent or family rights where relevant, and psychological consequence: preservation and respectful commemoration are not the same as deceptive resurrection, emotional exploitation, or

Conversational Agents for Promoting Mental Health and Well-Being," *npj Digital Medicine* 6 (2023): 236; Yuhao He et al., "Conversational Agent Interventions for Mental Health Problems: Systematic Review and Meta-analysis of Randomized Controlled Trials," *Journal of Medical Internet Research* 25 (2023): e43862; Bethanie Maples et al., "Loneliness and Suicide Mitigation for Students Using GPT3-Enabled Chatbots," *npj Mental Health Research* 3 (2024): 4; American Psychological Association, "Use of Generative AI Chatbots and Wellness Applications for Mental Health," APA Health Advisory, 2025; Laurie O. Campbell, Kathryn Babb, Glenn W. Lambie, and B. Grant Hayes, "An Examination of Generative AI Response to Suicide Inquires [sic]: Content Analysis," *JMIR Mental Health* 12 (2025): e73623, <https://doi.org/10.2196/73623>.

fantasy-based attachment.⁵⁸

5.8 Practical Fiqh Framework

A practical ruling framework may be summarized as follows:

1. Non-erotic, bounded, transparent AI companionship for a lawful benefit is permissible in origin when it does not involve deception, prohibited content, neglect of duties, privacy violation, or harmful dependency.
2. AI companionship becomes blameworthy or prohibited when it predictably cultivates emotional dependency, social withdrawal, neglect of family, or displacement of human obligations.
3. Romantic or erotic AI companionship is impermissible as it deliberately stimulates desire, imitates a spousal or illicit romantic relationship, involves sexualized exchange, encourages secrecy, undermines marital rights and marital health, exposes users to prohibited visual elements, or functions as a means of approaching zinā.
4. Use by minors requires especially strong restriction, and Muslim parents and institutions should not allow unsupervised access to companion systems designed for emotional attachment.

This framework preserves both caution and nuance. It avoids the error of treating AI companions as real spouses or legal persons. It also avoids the opposite error of assuming that artificiality eliminates moral consequence. The machine is not morally accountable, but the human heart is still affected; the AI has no desire, but the user may cultivate desire through it; the chatbot has no rights, but a spouse, child, parent, teacher, employer, or community may still be wronged by what the user does through it. In this sense, the fiqh of AI virtual relationships is ultimately not about the rights of machines. It is about the purification of human desire, the preservation of lawful relationships, and the refusal to let synthetic intimacy deform the moral architecture of Muslim and societal life.

⁵⁸ Jordanian General Iftaa' Department, "Ruling on Converting Photographs into Drawings Using Artificial Intelligence Technology," Fatwa no. 3989, July 24, 2025. The fatwa permits AI conversion of photographs into drawings when the content and form do not contradict Islamic rulings, values, and ethics, while prohibiting uses involving awrah, mockery, bullying, lying, forgery, fraud, or deceit.

6. Deepfakes and Digital Manipulation

Deepfakes and synthetic media present one of the clearest cases in which artificial intelligence threatens not merely individual conduct but the social conditions of truth, trust, reputation, and justice. For the purposes of this paper, a deepfake may be defined, in general terms, as AI-generated or AI-manipulated audio, image, or video content that realistically depicts a person saying, doing, appearing to do, or being present in something that did not occur in that way. The legal and ethical concern is not limited to technical falsity. Human communities rely on ordinary confidence in testimony, documentation, visual evidence, and personal reputation. When fabricated media becomes easy to create and difficult to distinguish from authentic media, harm is not confined to a single false item. It extends to the erosion of trust itself.

This section begins from a simple but important distinction. Not all synthetic media is unlawful. AI-generated images, dramatized reenactments, educational simulations, dubbing, accessibility tools, privacy-preserving synthetic data, and artistic work may be potentially permissible when they are clearly labeled, non-deceptive, lawful in content, respectful of rights, and free from reputational or financial harm. The problem arises when synthetic media is used to deceive, impersonate, defame, sexualize, manipulate, fabricate evidence, falsely attribute speech, or exploit another person's likeness, voice, or identity without consent. The ruling therefore depends not merely on the technology used, but on the legally relevant attributes of the act: deception, consent, content, harm, rights, public consequence, and the likelihood of confusion.

This discussion concerns synthetic media as it relates to deception, impersonation, consent, evidentiary trust, privacy, religious attribution, and harm. It does not attempt to adjudicate the broader juristic debate over image-making as such, which requires separate treatment. The point here is narrower: even where a form of image-making or synthetic production may be treated as permissible by some jurists or in some contexts, it may become prohibited through false attribution, deception, sexualization, reputational injury, violation of privacy, or use as fabricated evidence.

6.1 Truthfulness

The Shari'ah's concern for truthfulness is decisive. The Qur'an commands believers to verify serious reports before acting upon them, lest they harm people in ignorance and later regret what they have done.⁵⁹ It also warns against pursuing that of which one has no knowledge and condemns false witness.⁶⁰ The Prophet, peace be upon him, treated false testimony as among the gravest of major sins, repeating the warning until his Companions wished he would stop (out of concern for him).⁶¹ Deepfakes intensify the danger of false witness because they do not merely produce a false verbal claim; they manufacture an apparent sensory event. They give falsehood the emotional force of seeing and hearing.

This is why deepfakes should be analyzed not only under the general category of lying, but also under *shahādat al-zūr* (false testimony), *buhtān* (slandorous falsehood), *ifk* (fabricated calumny), *ghish* (deception or fraud), *tadlīs* (misrepresentation), impersonation, and violation of *amānah* (entrusted responsibility). In ordinary lying, the liar makes a false

⁵⁹ Qur'an 49:6.

⁶⁰ Qur'an 17:36; Qur'an 25:72.

⁶¹ *Ṣaḥīḥ al-Bukhārī*, no. 2654. The hadith identifies false witness among the gravest major sins and describes the Prophet, peace be upon him, repeatedly warning against it.

claim. In a deepfake, the producer may create a false “witness” by fabricating visual or auditory evidence. In commercial settings, this may amount to fraud or deceptive representation. In personal settings, it may amount to reputational injury. In political or communal settings, it may amount to manipulation of public judgment. In religious settings, it may amount to false attribution to a scholar, imam, reciter, or institution, thereby corrupting religious trust. In deepfake abuse, particularly while considering *ḥifẓ al-ʿird*, the victim’s “honor” is not an abstract social label. It includes reputation, dignity, identity, privacy, social standing, marital prospects, professional credibility, psychological safety, and communal belonging.

6.2 Deception and Fraud

The concept of *ghish* (deception or fraud) is also highly relevant. The Prophet, peace be upon him, said that whoever deceives is not from his way.⁶² Although the immediate hadith context concerns market deception, jurists understood the moral principle broadly: concealing the true nature of a thing in a way that misleads another party violates trust. Deepfakes are a technologically advanced form of *ghish* when they conceal artificial production and invite the viewer to believe that the depicted person truly spoke, acted, endorsed, confessed, threatened, recited, appeared, or consented. The deception is compounded when the media is shared by users who suspect or know it to be false, because dissemination becomes participation in the harm.

6.3 Sexual Deepfakes

The matter becomes even more severe when deepfakes involve sexual content or accusations of sexual misconduct. The Qur’an imposes severe consequences for those who accuse chaste persons without the required proof, and it condemns the spread of indecency among believers.⁶³ Classical *qadhḥ* (the legally punishable accusation of *zinā*, unlawful sexual intercourse, or denial of lineage) is a precise legal category and should not be applied carelessly to every manipulated image or defamatory insinuation. However, sexualized deepfakes involving identifiable persons may involve *qadhḥ*-like reputational destruction and, depending on the wording, context, and claim being communicated, may enter *qadhḥ* itself. Even where the formal *ḥadd* category is not established, the act remains a grave violation of *ʿird* (honor/reputation), privacy, and dignity, and it may warrant *taʿzīr* (discretionary punishment) and compensation where judicial authority exists.

The classical apparatus for *qadhḥ* sharpens this analysis. The jurists distinguished the explicit accusation (*ṣarīḥ*) from insinuation (*taʿrīḍ*), and they disagreed over the latter: the Mālikīs imposed the *ḥadd* for insinuation that is understood by its audience as an accusation, following the precedent of ‘Umar ibn al-Khaṭṭāb, who applied the punishment to a man whose pointed denial about his own parents was understood by all present as an accusation against another’s, while the majority confined the *ḥadd* to explicit wording and assigned insinuation to *taʿzīr*, invoking the principle that prescribed punishments are

⁶² *Ṣaḥīḥ Muslim*, no. 102.

⁶³ Qur’an 24:4; Qur’an 24:19; Qur’an 24:23; *Ṣaḥīḥ al-Bukhārī*, no. 6857, listing accusations of chaste believing women among the destructive sins.

averted by doubt.⁶⁴

A sexualized deepfake implying *zinā* occupies an interesting position within this framework. It utters no words at all, which on the majority's formal analysis keeps it outside the ḥadd; yet it communicates the accusation more forcefully than any verbal *ta'rīd* in history, because it does not merely imply that the victim committed an act but manufactures the apparent sensory proof of it. On the Mālikī criterion of what the audience necessarily understands, the fabricated depiction is the most unambiguous accusation conceivable. This paper does not presume to resolve the ḥadd question, which belongs to judicial authority with its evidentiary standards; the practical conclusions hold on every madhhab's analysis.⁶⁵ At minimum the act is a grave *ta'zīr* offense combining the moral substance of *qadhf* with the fabrication of false evidence, warranting punishment, compensation for demonstrable harm, removal of the content, and public restoration of the victim's reputation. And in one respect it exceeds verbal *qadhf* in gravity: the verbal accuser makes a claim that the Shari'ah then tests against *bayyinah*, while the deepfake producer counterfeits the *bayyinah* itself, attacking not only the victim's honor but the community's machinery for distinguishing truth from falsehood.

These points require particular emphasis because nonconsensual intimate deepfakes have become one of the most widespread and harmful forms of synthetic media abuse. Contemporary scholarship describes nonconsensual intimate deepfakes as a form of image-based sexual abuse and notes that they disproportionately harm women and girls.⁶⁶ This modern empirical observation maps directly onto the Shari'ah's protection of 'ird, chastity, privacy, and human dignity. A person need not physically touch the victim to commit a severe violation. To fabricate a sexual image of a chaste person, circulate it, threaten to circulate it, or derive enjoyment, profit, or leverage from it is a serious act of oppression, falsehood, and reputational violence.

6.4 Deepfakes and Epistemology

Deepfakes also require an epistemological analysis. Islamic law does not treat every claim as proof. The claimant must bring evidence, and mere assertion is insufficient to violate another person's life, property, marriage, or reputation. The Prophet, peace be upon him, stated that if people were given according to their claims, some would claim the lives

⁶⁴ For the 'Umar precedent, see Mālik, *Al-Muwatta'*, Kitāb al-Ḥudūd, Bāb al-Ḥadd fi al-Qadhf wa-l-Nafy wa-l-Ta'rīd, no. 1521. On the maxim that ḥudūd are averted by doubts, see al-Suyūṭī, *Al-Ashbāh wa-l-Nazā'ir fi Qawā'id wa-Furū'* Fiqh al-Shāfi'iyyah (Beirut: Dār al-Kutub al-'Ilmiyyah, 1403/1983), 122.

⁶⁵ Ibn Rushd, *Bidāyat al-Mujtahid wa-Nihāyat al-Muqtaṣid*, Kitāb al-Qadhf, under "al-qadhf alladhī yajibū bihi al-ḥadd," 2:440 in the al-Bābī al-Ḥalabī edition. Ibn Rushd states that they agreed that explicit qadhf incurs the ḥadd, then disagreed over ta'rīd. He reports that al-Shāfi'ī, Abū Ḥanīfah, al-Thawrī, and Ibn Abī Laylā held there is no ḥadd for ta'rīd, though Abū Ḥanīfah and al-Shāfi'ī held ta'zīr applies. Mālik and his companions held that ḥadd applies for ta'rīd, and Ibn Rushd connects the issue to the incident in 'Umar's time. He also states the majority's reasoning: the ambiguity in figurative wording is a shubha, and ḥudūd are averted by doubts.

⁶⁶ Beatriz Kira, "When Non-Consensual Intimate Deepfakes Go Viral," *Computer Law & Security Review* 54 (2024): 106024. Kira describes nonconsensual intimate deepfakes as a form of image-based sexual abuse and notes their disproportionate impact on women and girls.

and wealth of others, but the oath is upon the defendant.⁶⁷ This principle is deeply relevant in an age of synthetic media. Videos, screenshots, audio recordings, images, and chat logs may function as *qarā'in* or supporting indicators, but they cannot be treated with naïve certainty when manipulation is technically accessible and increasingly realistic.⁶⁸ Deepfakes do not abolish digital evidence, but they demand a more disciplined evidentiary ethic.

The juristic concern here is not only that fake media may be believed. It is also that real media may be dismissed as fake. Chesney and Citron identify this as the "liar's dividend": as people become aware that convincing audio and video can be fabricated, wrongdoers gain an opportunity to deny authentic evidence by claiming it is artificial.⁶⁹ Political life has already begun to exhibit this danger, and no doubt it will continue to play a role in shaping narratives and political agendas.⁷⁰

The point is not the partisan identity of any example that can be cited neutrally, but the evidentiary pattern: once the public knows that synthetic media exists, a public figure may exploit that uncertainty to cast doubt on inconvenient reality. Political scientists have likewise argued that strategic claims of "fake news" or "deepfake" can help politicians preserve support after scandal by invoking informational uncertainty or rallying supporters against perceived hostile media. This phenomenon has a direct parallel in the erosion of *yaqīn* (certainty) and *ghalabat al-ẓann* (preponderant or dominant probability) in communal judgment. When the public can no longer distinguish truth from fabrication, the wrongdoer benefits twice: first from the ability to fabricate falsehood, and second from the ability to deny truth. This constitutes a profound harm to *bayyinah* (clear proof), *shahādah* (testimony), and public trust.

6.5 Guidelines

For imams, community leaders, arbitrators, and Muslim institutions, this means that digital media must be handled with verification, restraint, and procedural caution. Qur'an 49:6 already establishes the obligation of verification when reports may harm others. In a deepfake environment, verification includes assessing provenance, context, source reliability, metadata where available, corroborating evidence, expert analysis when necessary, chain of custody, and the possibility of manipulation. It also includes refraining

⁶⁷ *Ṣaḥīḥ Muslim*, no. 1711a.

⁶⁸ Synthetic or digital media, by itself, should not be treated as proof for establishing hudud in applicable contexts; the strict evidentiary standards of the Sharī'ah remain intact, while such media may be investigated as *qarā'in* in matters of ta'zīr, civil harm, or protective action according to authority and context.

⁶⁹ Robert Chesney and Danielle Keats Citron, "Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security," *California Law Review* 107, no. 6 (2019): 1753–820; Kaylyn Jackson Schiff, Daniel S. Schiff, and Natália S. Bueno, "The Liar's Dividend: Can Politicians Claim Misinformation to Evade Accountability?," *American Political Science Review* 119, no. 1 (2025): 71–90, <https://doi.org/10.1017/S0003055423001454>; Josh A. Goldstein and Andrew Lohn, "Deepfakes, Elections, and Shrinking the Liar's Dividend," Brennan Center for Justice, January 23, 2024.

⁷⁰ For example, during the 2024 United States presidential election, Donald Trump falsely claimed that a real photograph of a Kamala Harris campaign crowd in Michigan was AI-generated to exaggerate attendance; Reuters reported that the image was authentic, that the Harris campaign confirmed it was taken by a staff member, and that expert review and other contemporaneous photos and videos supported the reality of the crowd. See Reuters, "Photo from Harris-Walz Rally in Michigan Was Not AI-Generated," August 12, 2024.

from public commentary when verification is incomplete and reputational harm would be severe. A masjid board, school, family mediator, or community leader should not act on viral media alone when the consequence may involve dismissal, divorce, public shaming, police reporting, or communal scandal.

The principle *al-yaqīn la yazūl bi-al-shakk* is also relevant, though it must be applied carefully. Certainty is not removed by mere doubt, but deepfake possibility does not render all evidence doubtful. Rather, the level of verification required rises according to the seriousness of the consequence and the plausibility of manipulation. A casual meme does not require the same scrutiny as an allegation of zinā, abuse, fraud, or any criminal conduct. The greater the harm attached to accepting the media, the stronger the need for verification. Conversely, the mere possibility of deepfake manipulation should not become a shield for every wrongdoer. The jurist must avoid both naïve acceptance and cynical dismissal.

6.6 Religious Deepfakes

Religious deepfakes deserve distinct treatment. It is now technically possible to imitate a scholar's voice, generate a fake video of an imam endorsing a product or political cause, fabricate a clip of a reciter, produce synthetic khutbahs attributed to living or deceased scholars, or generate false fatwas using a scholar's likeness. This category is especially dangerous because it combines false attribution, religious authority, and public trust. To fabricate religious speech and attribute it to a scholar is not merely ordinary impersonation; it is a violation of *amānah* in the transmission of religion. It may mislead worshippers, damage scholarly reputations, create false communal conflict, or spread rulings that people then act upon in worship, marriage, finance, or family life.

The Sharī'ah's warning against lying about the Prophet ﷺ is uniquely severe and cannot be analogized simplistically to lying about others in every respect.⁷¹ Yet the underlying concern for truthful religious transmission is central to Islam. Classical scholarship treated false attribution to the Prophet ﷺ with exceptional gravity because his speech conveys religion, law, and guidance, so fabrication in his name corrupts the transmission of dīn itself. Scholars of hadith therefore built entire disciplines around attribution, verification, narrator reliability, and the protection of the Sunnah from forgery. A religious deepfake violates that ethic of *isnād* and *amānah* by manufacturing the appearance of transmission without its reality. Muslim institutions should therefore adopt a strict presumption against the use of synthetic likenesses or voices of scholars, reciters, or religious figures without explicit consent and clear disclosure.

The same applies to fake endorsements and institutional communications. If a donor, student, congregant, or consumer is led to believe that an imam, scholar, masjid, charity, or Islamic seminary endorsed a product, candidate, fatwa, fundraiser, or position when it did not, the act combines *ghish*, false attribution, and possible financial fraud. Where money is collected, reputations are altered, or religious trust is exploited, the matter may move beyond personal sin into liability, restitution, and public correction. If false media spreads, those who knowingly shared it bear responsibility proportionate to their knowledge, intent, reach, and role in harm.

6.7 Consent Over One's Likeness

A further issue is consent over one's image, voice, likeness, and biometric identity.

⁷¹ *ṣaḥīḥ al-Bukhārī*, no. 106, *ṣaḥīḥ Muslim*, no. 3.

Classical fiqh did not address AI-generated voice cloning or facial synthesis in these exact forms, but it did establish principles from which their ruling may be derived. By *taḥqīq al-manāṭ*, the effective causes that trigger protection in the classical law are generally present here: the safeguarding of *ʿird*, the sanctity of private life,⁷² the prohibition of wrongful harm,⁷³ and the prevention of deception and exploitation through another person's identity. The Sharī'ah forbids spying, requires permission before entering private spaces,⁷⁴ condemns harming believers without right, and treats a Muslim's honor as inviolable.⁷⁵

It also recognizes, through valid *ʿurf* and contemporary juristic development, that certain incorporeal interests may become legally cognizable and protected rather than treated as ownerless raw material.⁷⁶ On that basis, a person's recognizable likeness and voice should not be treated as public material open to extraction, manipulation, or monetization without consent. They are extensions of protected personal identity in which interests of privacy, reputation, dignity, and, in some circumstances, compensable benefit converge. This conclusion is further reinforced by *maṣlaḥah* and *sadd al-dharāʾi*: permitting unrestricted synthetic appropriation of identity would invite deception, exploitation, reputational injury, and the erosion of trust in testimony and attribution. The Qur'anic affirmation that Allah has honored the children of Adam supports this analysis,⁷⁷ but the case rests more broadly on the Sharī'ah's protection of persons from humiliation, intrusion, and harm.

At the same time, a nuanced fiqh framework must recognize conditions under which synthetic media may be permissible. These include cases where the content is lawful, the purpose is legitimate, the synthetic nature is clearly disclosed, no identifiable person is impersonated without consent, no reputation is harmed, no private information is exposed, no sexualized or degrading depiction is created, and no audience is reasonably misled. Educational reenactments, historical simulations, clean media that cannot reasonably be mistaken for reality, accessibility dubbing, anonymized training data, and forensic training

⁷² See Qur'an 24:27; 49:12. These verses ground the Sharī'ah's concern for private space, permission, and the prohibition of intrusive surveillance, all of which are relevant by extension to unauthorized capture, imitation, or synthetic reproduction of a person's identity-bearing features

⁷³ On the maxim *lā ḍarar wa lā ḍirār* ("There should be neither harming nor reciprocating harm"), see al-Nawawī, *al-Arbaʿūn*, no. 32. This maxim supports extending legal protection to new technologies that inflict identifiable personal, reputational, sexual, or financial harm.

⁷⁴ See *Ṣaḥīḥ Muslim*, no. 2158a; *Ṣaḥīḥ al-Bukhārī*, no. 6901. The Prophet ﷺ treated unlawful visual intrusion into private space with exceptional seriousness, even stating that one who peeps into a house without consent may have his eye struck out, and explaining that permission to enter was legislated precisely to prevent unauthorized looking.

⁷⁵ See Qur'an 33:58; *Ṣaḥīḥ Muslim*, no. 2564a ("All things of a Muslim are inviolable for his brother in faith: his blood, his wealth and his honour"). Together, these sources support the protection of persons from unjust injury to honor and reputation.

⁷⁶ See International Islamic Fiqh Academy, Resolution No. 43 (5/5) on Moral Rights, which states that trade names, trademarks, literary production, inventions, and discoveries are rights belonging to their holders, recognized by Sharī'ah, and not permissible to violate. While a person's likeness is not identical to each of these categories, the resolution is significant because it confirms that modern incorporeal interests may be legally cognizable in Sharī'ah rather than dismissed as ownerless abstractions.

⁷⁷ Qur'an 17:70.

may fall into this category. The permissibility is not rooted in the artificiality of the media alone, but in the absence of deception, rights violation, and harm.

6.8 Modern governance frameworks

Modern governance frameworks are converging on similar concerns, especially around labeling and provenance. The European Union's AI Act includes transparency obligations for certain AI systems, including informing users when they interact with AI and requiring disclosure or marking of AI-generated or manipulated content, including deepfakes, subject to specified conditions and exceptions.⁷⁸ NIST's synthetic content risk report likewise discusses provenance data, watermarking, and detection as means to help establish the authenticity, integrity, and credibility of digital content.⁷⁹ These measures are useful but imperfect; provenance metadata, watermarking, and detection can be degraded, removed, bypassed, or outpaced by adversarial generation. Furthermore, provenance standards, watermarking, disclosure labels, and content credentials are practical attempts to operationalize amanah, even though they are not foolproof and must be supplemented by verification and accountability. These measures do not fully define the Sharī'ah ruling, but they support a practical conclusion: transparency, provenance, labeling, and verification are necessary safeguards for preserving trust in an age of synthetic media.

Responsibility for deepfake harm is layered and should be assessed according to causation, intent, knowledge, control, and negligence. The creator bears the clearest responsibility when he intentionally fabricates deceptive or harmful content. The one who commissions, directs, or knowingly deploys it shares responsibility according to his role in bringing about the harm. The one who knowingly shares it participates in dissemination, while the one who shares it without verification may also bear blame where serious harm was reasonably apparent and precaution was neglected.⁸⁰ Platforms and distributors bear increasing responsibility when they knowingly facilitate, monetize, recommend, or fail to remove harmful synthetic content after credible notice, especially in cases involving sexual abuse, fraud, incitement, or severe reputational injury.

Upstream developers and infrastructure providers may likewise bear responsibility when their contribution is not merely remote or general-purpose,⁸¹ but is intentionally

⁷⁸ Regulation (EU) 2024/1689, Artificial Intelligence Act, art. 50. The European Commission summarizes the Act's transparency obligations as including marking AI-generated content and disclosing the artificial nature of images, audio, deepfakes, and text.

⁷⁹ National Institute of Standards and Technology, Reducing Risks Posed by Synthetic Content: An Overview of Technical Approaches to Digital Content Transparency, NIST AI 100-4 (Gaithersburg, MD: NIST, 2024). The report discusses provenance information, watermarking, and synthetic content detection as methods for supporting authenticity and integrity.

⁸⁰ A person who knowingly spreads harmful synthetic content falls under the prohibition of cooperation in sin and transgression. A person who spreads it unknowingly is not identical to the knowing disseminator, but he is still governed by the Qur'anic obligation of verification and the prohibition of speaking without knowledge, especially where reputational, sexual, or financial harm is at stake.

⁸¹ Contemporary juristic guidance on AI also supports differentiating responsibility across the development and deployment stack. The International Islamic Fiqh Academy's 2025 AI resolution states that the default rule for creating and using AI is permissibility, but it

directed toward abuse, predominantly optimized for unlawful impersonation or exploitation, or left without meaningful safeguards despite repeated and serious evidence of misuse.⁸² These layers correspond, in contemporary application, to familiar fiqh distinctions concerning *mubāsharah*, *tasabbub*, *ta'āwun 'alā al-ithm*, *ta'addī*, *tafrīt*, and *ḍamān*.

6.9 Psychological Impact of Deepfake Abuse

Deepfake abuse also requires attention to harms that are difficult to measure but morally grave. A synthetic video, forged voice, or sexualized image may not physically touch the victim, yet it can damage *'ird*, expose private dignity, rupture family trust, destroy professional credibility, produce social stigma, and cause severe psychological distress. A recent Islamic-ethics study on deepfake abuse therefore places *ḥifẓ al-'ird* at the center of the maqāṣid analysis and argues that reputational destruction, blackmail, political manipulation, and nonconsensual sexual fabrication should be treated as violations of protected human dignity rather than as merely technical misuse of media.⁸³ This is an important corrective to legal frameworks that often focus most easily on material losses while struggling to quantify reputational, psychological, and communal harm. From a Shari'ah perspective, the moral seriousness of deepfakes does not depend only on financial damage or physical injury. The destruction of honor, the spread of false attribution, and the forced exposure of a person to public humiliation are themselves forms of harm that must be prevented, removed, and remedied. Thus, in cases of deepfake abuse, a Shari'ah-informed response should not stop at condemning the perpetrator. It should also seek *izālat al-ḍarar* through removal of harmful content, correction of false narratives, restoration of reputation, pastoral support, psychosocial care, and institutional assistance for victims.

6.10 Civic Engagement

A Shari'ah-informed response to deepfakes should also include civic advocacy. Muslims should not treat AI regulation as a distant technical matter reserved for engineers, corporations, or national-security officials; it is now a question of public trust, dignity,

conditions that on the legitimacy of the purpose, use, financing, and outcome, as well as on protecting information and safeguarding public and private rights, and observing honesty, documentation, and transparency in use. This supports assigning greater responsibility where a system is specifically directed, financed, or deployed toward harmful impersonation, rather than treating all upstream contributors identically. International Islamic Fiqh Academy, Resolution No. 258 (3/26), "Artificial Intelligence: Its Rulings, Guidelines, and Ethics."

⁸² For this reason, modern references to "foreseeable misuse" should be framed cautiously in Shari'ah analysis. The more precise inquiry is not whether misuse was merely imaginable, but whether the actor intentionally facilitated it, strongly expected it, retained meaningful control over prevention, or failed in safeguards after serious evidence of recurring abuse. This is closer to the juristic logic of *ta'addī* and *tafrīt* than to an unlimited modern tort concept of foreseeability.

⁸³ Wisnu Uriawan et al., "Preventing AI Deepfake Abuse: An Islamic Ethics Framework," 2025, which argues that deepfake abuse implicates *ḥifẓ al-'ird* and *ḥifẓ al-naḥs* through reputational damage, psychological trauma, social exclusion, and economic harm; see also Luisa Verdoliva, "Media Forensics and DeepFakes: An Overview," IEEE Journal of Selected Topics in Signal Processing 14, no. 5 (2020): 910–32; Robert Chesney and Danielle Keats Citron, "Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security," California Law Review 107, no. 6 (2019): 1753–820.

privacy, election integrity, family protection, and protection from reputational and sexual exploitation. Recent legal developments show both urgency and incompleteness. The European Union's AI Act includes transparency obligations for certain AI-generated and manipulated content, including deepfakes, while U.S. state legislatures have introduced a rapidly growing number of AI-related bills and the federal Take It Down Act criminalizes knowingly publishing or threatening to publish nonconsensual intimate images, including AI-generated deepfakes, while requiring covered platforms to remove such material after proper notice. These are important but limited steps. They do not, by themselves, resolve broader questions of platform incentives, privacy, recommender-system harm, corporate concentration, political manipulation, religious deepfakes, biased outputs, or the use of AI infrastructure for state and military purposes.

In the United States, the need for moral vigilance is heightened by the close public alignment between major AI firms and state power; such developments show that AI governance is being shaped within a political economy of corporate power, national competition, deregulation, defense interests, and public-private partnership.⁸⁴ Muslims should therefore advocate for stronger safeguards at local, state, and federal levels, and press candidates for office to treat AI harms not merely as innovation policy but as matters of justice, dignity, privacy, truthfulness, family protection, and public amanah. The Islamic contribution here is not sectarian self-interest; it is a moral vocabulary that can benefit society by insisting that technological power must be accountable to the protection of human beings rather than merely to speed, profit, or geopolitical dominance.

6.11 Practical fiqh framework

A practical ruling framework may therefore be summarized as follows:

1. Creating or sharing synthetic media is permissible in origin when the content is lawful, the purpose is lawful, the synthetic nature is clear, consent is obtained where a person's likeness or voice is used, and no deception or harm is reasonably expected.
2. It is prohibited to create, request, possess for harmful use, share, or benefit from deepfakes that impersonate real persons without consent, fabricate speech or conduct, deceive audiences, damage reputation, create sexualized depictions, produce false evidence, manipulate religious authority, or facilitate fraud.
3. Sexual deepfakes are prohibited and especially grave, as they violate dignity, privacy, chastity, and reputation.
4. Religious deepfakes require heightened prohibition where they falsely

⁸⁴ For example, the Trump administration revoked the prior executive order requiring certain high-risk AI developers to share safety-test results with the government, announced the \$500 billion Stargate infrastructure project at the White House with OpenAI's Sam Altman and other technology executives present, and OpenAI later received a \$200 million U.S. Defense Department contract for national-security AI capabilities. See Barbara Ortutay, "President Trump Signs Take It Down Act, Addressing Nonconsensual Deepfakes. What Is It?," Associated Press, May 19, 2025; Kanishka Singh, "OpenAI Wins \$200 Million US Defense Contract," Reuters, June 16, 2025; Krystal Hu and Stephen Nellis, "Trump Announces Private-Sector \$500 Billion Investment in AI Infrastructure," Reuters, January 21, 2025; Associated Press, "Trump Highlights Partnership Investing \$500 Billion in AI," January 21, 2025; "Trump Revokes Biden Executive Order on Addressing AI Risks," Reuters, January 21, 2025; White House, "Winning the AI Race: America's AI Action Plan," July 2025.

attribute speech, recitation, endorsement, or fatwa to scholars, imams, institutions, or reciters.

5. Digital media used as evidence must be verified according to the seriousness of the claim and the harm that would follow from accepting it.

The deeper danger of deepfakes is that they undermine the moral ecology of truth. Communities cannot function when sight and sound are no longer presumptively trustworthy, when accusations spread faster than verification, when reputations can be destroyed by fabrication, and when authentic wrongdoing can be dismissed as artificial. Islamic law is not unprepared for this challenge. Its insistence on *tabayyun*, *bayyinah*, *shahādah*, *amanah*, protection of *'ird*, prohibition of *qadhf*, condemnation of false witness, and restraint from spreading indecency provides a powerful framework for synthetic media. The challenge is to apply these principles before deception becomes normalized and before communities learn, too late, that the loss of trust is far easier to cause than to repair.

7. Daily AI Conduct and Moral Harm

Daily-use AI affects Muslim conduct not only through exceptional cases such as deepfakes or virtual romantic partners, but through ordinary patterns of writing, studying, working, communicating, consuming media, seeking validation, and presenting oneself to others. This section concerns the moral harms that arise when AI facilitates deception, misrepresentation, dependency, exposure to prohibited content, or evasion of responsibility. The central question is not whether AI assistance is inherently unlawful. It is not. The question is when assistance becomes deception, when convenience becomes negligence, when personalization becomes manipulation, and when repeated use weakens amanah, truthfulness, restraint, and moral agency.

The starting point is that AI may be used for many lawful daily purposes. A Muslim may use AI to organize notes, improve grammar, translate non-sacred text, brainstorm lawful ideas, summarize permissible material, improve accessibility, schedule tasks, learn technical skills, or receive tutoring support. Such uses remain permissible in origin when they do not involve deception, unlawful content, violation of rights, breach of confidentiality, or harmful reliance. This is consistent with the broader juristic treatment of worldly tools and customary means: the tool is not condemned merely because it is new, but the ruling changes according to use, consequence, and rights implicated.

7.1 Deception of Competence

The first major category of concern is deception of competence. AI becomes impermissible when a person uses it to represent knowledge, skill, authorship, judgment, or effort that he or she does not actually possess in a context where those qualities are being assessed or relied upon. A student who submits AI-generated work as personal writing where personal authorship is required has not merely used a tool; he has misrepresented competence. A job applicant who uses AI to fabricate experience, exaggerate achievements, or produce false professional evidence has violated amanah. A professional who relies on AI to produce work beyond his competence, while presenting it as his own qualified judgment, may deceive clients, employers, patients, congregants, or the public.

This is why AI-related academic misconduct should not be reduced to the phrase “plagiarism” alone. Plagiarism is one form of academic dishonesty, but the deeper fiqh category is *ghish* and *khiyanat al-amanah*: deception in a trust. The Prophet, peace be upon him, said, “[Whoever deceives is not from me.](#)”⁸⁵ The Qur’an commands believers to fulfill trusts and not betray Allah, the Messenger, or the trusts placed in them.⁸⁶ In educational settings, an assessment is not merely a writing product; it is often a certification of learning, discipline, reasoning, memory, skill, and readiness for future responsibility. If AI use conceals the absence of those qualities, the harm may extend beyond the classroom to professional qualification, public safety, employment fairness, and institutional trust.

Contemporary education literature supports this distinction. Kasneci and colleagues note that large language models offer significant educational opportunities, including personalized learning support, content generation, and feedback, but also require learners

⁸⁵ *Ṣaḥīḥ Muslim*, no. 102.

⁸⁶ Qur’an 4:58; Qur’an 8:27.

and teachers to develop literacy regarding their limitations and unexpected brittleness.⁸⁷ UNESCO's guidance on generative AI in education likewise emphasizes human agency, data privacy, age-appropriate use, and institutional governance rather than unregulated adoption.⁸⁸ A systematic review by Balalle similarly concludes that AI creates both opportunities and serious academic integrity challenges, including risks to the credibility of qualifications.⁸⁹ These sources help clarify the *wāqī'*: the problem is not AI-assisted learning as such, but undisclosed or unauthorized AI substitution where learning, authorship, or competence is being evaluated.

A practical distinction is therefore necessary. AI may be used as an editor, tutor, brainstorming partner, translation aid, or accessibility support when the institution permits it, the user remains intellectually responsible, and the use does not misrepresent what is being assessed. It becomes prohibited when the user submits AI output as his own work in a context requiring personal effort, uses AI to evade learning, fabricates sources, conceals unauthorized assistance, or gains an unfair advantage over others. Between these cases are gray areas requiring institutional clarity. Schools, seminaries, universities, and employers should define acceptable and unacceptable AI use in advance, since ambiguity can create both genuine confusion and deliberate exploitation.

Where a school, employer, publisher, seminary, or professional body clearly discloses conditions governing AI use, compliance with those conditions becomes part of the relevant condition, recognized *urf*, and *amanah*. Violating such policies is therefore not merely a breach of administrative preference; it may become deception and betrayal of trust when grades, certification, employment, publication, religious instruction, or public reliance are affected. This is especially important because many AI-related harms arise not from the tool itself, but from concealing its use in a context where personal effort, authorship, competence, or judgment is being evaluated.

The duty of truthfulness in this domain runs in both directions, and institutions must hear the second direction clearly. The same evidentiary discipline this paper demands before acting on a deepfake applies before accusing a student or employee of concealed AI use. Automated AI-text detectors are statistical classifiers, not witnesses: they produce false positives at meaningful rates, their judgments cannot be audited or cross-examined, and published research has found them disproportionately likely to flag the genuine writing of non-native English speakers, while at least one major developer withdrew its own

⁸⁷ Enkelejda Kasneci et al., "ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education," *Learning and Individual Differences* 103 (2023): 102274. Kasneci and colleagues discuss both educational opportunities and the need for teachers and learners to develop competencies and literacies concerning the limitations and brittleness of large language models.

⁸⁸ Fengchun Miao and Wayne Holmes, *Guidance for Generative AI in Education and Research* (Paris: UNESCO, 2023). UNESCO's guidance emphasizes human agency, protection of data privacy, age-appropriate approaches, and ethical validation in education and research.

⁸⁹ Himendra Balalle and Sachini Pannilage, "Reassessing Academic Integrity in the Age of AI: A Systematic Literature Review on AI and Academic Integrity," *Social Sciences & Humanities Open* 11 (2025): 101299, <https://doi.org/10.1016/j.ssaho.2025.101299>. Balalle and Pannilage's systematic review examines the impact of AI on academic integrity and emphasizes the need to balance AI's benefits with the preservation of ethical standards and the credibility of qualifications.

detection tool over its unreliability.⁹⁰

The same caution applies to false negatives. A detector's failure to identify AI-generated text does not establish that no unauthorized AI use occurred, just as a positive result does not establish that it did. Whether positive or negative, the classifier's output is at most a *qarīnah* to be weighed alongside other evidence; it is neither a witness nor independently dispositive in either direction. Institutions should therefore anticipate evidentiary needs when designing assessments and integrity policies, rather than attempting to reconstruct a student's or employee's process only after suspicion arises. Draft histories, version records, research notes, staged submissions, documented feedback, and, where appropriate, an oral defense can provide more reliable evidence of authorship, competence, and intellectual process. The corroboration required for a just determination will often be available only if institutional policy and assessment design have preserved it in advance.

An accusation of *ghish* is a claim against a person's *ʿird*, religious standing, and livelihood, and the Shari'ah's rule is that the burden of proof rests on the claimant, that the accused retains the presumption of innocence, and that the required strength of evidence rises with the severity of the consequence. A detector score is therefore at most a *qarīnah* that justifies inquiry; it is never by itself bayyinah that justifies sanction. An institution that expels, fails, fires, or publicly shames a person on the basis of an opaque classifier's output, without corroboration through drafts, process records, oral examination, or the person's own response, has not enforced amānah but violated it, replacing one form of automated injustice with another. Tabayyun is owed to the accused student no less than to the victim of a fabricated video. Muslim institutions should write this into policy: detector output may open a conversation, never close a case; the accused must be heard; corroborating evidence of process must be sought; and the gravity of the sanction must be matched by the gravity of the proof.⁹¹

7.2 Deception in Professional, Commercial, and Legal Communication

The second category is deception in professional, commercial, and legal communication. AI can generate resumes, business proposals, contracts, customer messages, fundraising appeals, legal drafts, medical summaries, religious content, reviews, and public statements. Lawful assistance in wording is not inherently problematic. The problem arises when AI is used to fabricate credentials, exaggerate expertise, forge testimonials, produce fake reviews, impersonate another person, conceal material facts, or generate confident claims that the user has not verified. The Qur'an's condemnation of fraud in measure and weight expresses a broader moral principle: one must not take from others on the basis of distorted representation.⁹² In modern contexts, competence, authorship, data, credentials, and professional judgment are also "measures" by which people transact and trust.

AI makes such deception easier because it lowers the effort required to appear

⁹⁰ Weixin Liang et al., "GPT Detectors Are Biased Against Non-Native English Writers," *Patterns* 4, no. 7 (2023): 100779, <https://doi.org/10.1016/j.patter.2023.100779>; OpenAI, "New AI Classifier for Indicating AI-Written Text," January 31, 2023, updated July 2023 to note that the classifier was discontinued due to its low rate of accuracy.

⁹¹ *Ṣaḥīḥ Muslim*, no. 1711a.

⁹² Qur'an 83:1-3.

informed, sincere, polished, experienced, or authoritative. A weak student can appear learned, an unqualified applicant can appear competent, a negligent employee can appear diligent, a manipulative suitor can appear eloquent, and a dishonest seller can generate persuasive claims at scale. The sin is not eloquence itself, nor the use of assistance in expression. The sin is the deliberate creation of a false impression in a context where the other party has a right to truthful disclosure. This falls under *ghish*, *tadlis*, false testimony where relevant, and betrayal of *amanah*.

7.3 Impersonation and False Relational Communication

The third category is impersonation and false relational communication. A person may use AI to write messages that are more emotionally persuasive than his actual intention, generate apologies he does not mean, create romantic communication to manipulate another person, or maintain multiple identities. Again, assistance is not automatically impermissible. A person may seek help expressing a sincere apology or translating a message. But if AI is used to produce emotional manipulation, conceal identity, sustain a false persona, deceive a prospective spouse, or imitate another person's voice or writing style without consent, the ruling changes. The Shari'ah's concern is not merely the mechanics of authorship, but the violation of truth, consent, and trust.

7.4 Algorithmic Exposure to Prohibited or Harmful Content

The fourth category is algorithmic exposure to prohibited or harmful content. AI systems increasingly mediate what users see, hear, read, and desire. Recommender systems, conversational assistants, search tools, and personalized feeds can expose users to sexual content, gambling-like engagement, extremist material, self-harm content, ideological rage, consumer manipulation, or spiritually corrosive entertainment. The user remains responsible for voluntary pursuit of prohibited content, but the platform or designer may also bear responsibility when harm is foreseeable, monetized, and amplified. The Qur'an prohibits cooperation in sin and transgression, while commanding cooperation in righteousness and piety.⁹³ This principle is important in evaluating not only the end-user but also the systems that profit from user weakness.

Modern research on recommender systems is nuanced but relevant. A systematic review by Yesilada and Lewandowsky examined whether YouTube's recommender system facilitates pathways to problematic content such as extremist material and found a mixed but serious body of evidence requiring careful interpretation.⁹⁴ Whittaker and colleagues similarly examined recommendation systems across platforms and found that recommendations can amplify problematic content in certain conditions, while also

⁹³ Qur'an 5:2.

⁹⁴ M. Yesilada and S. Lewandowsky, "Systematic Review: YouTube Recommendations and Problematic Content," Internet Policy Review 11, no. 1 (2022). The review examines whether YouTube's recommender system facilitates pathways to problematic content such as extremist or radicalizing material.

cautioning against overly simplistic accounts of algorithmic radicalization.⁹⁵ The fiqh implication is precisely this nuance: one should not claim that every algorithm necessarily radicalizes users, but neither should one ignore the moral relevance of systems that predictably amplify harmful material for engagement.

7.5 Addictive or Compulsive Engagement

The fifth category is addictive or compulsive engagement. The term “addiction” must be used carefully. AI addiction is not yet a settled diagnostic category in the way some related disorders are described in clinical classifications. The World Health Organization’s ICD-11 defines gaming disorder as a pattern of gaming behavior marked by impaired control, increasing priority over other activities, and continuation or escalation despite negative consequences.⁹⁶ This does not prove that AI use is medically identical to gaming disorder, intoxicants, or gambling. It does, however, provide a useful analogy for identifying when a technology becomes morally dangerous: loss of control, domination of attention, neglect of obligations, and continued use despite harm.

The Qur’an’s treatment of *khamr* and *maysir* is instructive as a model of balancing benefit and harm. It acknowledges that wine and gambling contain some benefit, while declaring that their sin is greater than their benefit, and later prohibits them as abominations from the work of Satan.⁹⁷ The analogy to AI must be precise. AI is not *khamr* and is not *maysir* in itself. But certain AI-enabled systems may resemble morally relevant features of gambling or intoxication when they exploit reward loops, impair judgment, encourage compulsive return, create artificial hope, or monetize loss of self-control. This is especially relevant to systems designed around endless novelty, flattery, sexual gratification, rage engagement, or personalized emotional reinforcement.

Recent AI behavior also illustrates the danger of sycophancy. OpenAI publicly acknowledged in 2025 that a GPT-4o update had made the model overly flattering and agreeable, including in ways that could validate doubts, fuel anger, urge impulsive actions, or reinforce negative emotions; the company stated that such behavior can raise safety concerns involving mental health, emotional over-reliance, and risky behavior.⁹⁸ Another recent study found that AI systems affirmed users’ actions about 49 percent more often than humans, including in scenarios involving deception, illegal conduct, or socially harmful behavior; the authors found that sycophantic responses could increase users’ conviction

⁹⁵ Joe Whittaker, Seán Looney, Alastair Reed, and Fabio Votta, “Recommender Systems and the Amplification of Extremist Content,” *Internet Policy Review* 10, no. 2 (2021). The authors provide empirical analysis of recommendation systems across platforms in relation to far-right content and discuss policy concerns about amplification and filter bubbles.

⁹⁶ World Health Organization, “Gaming Disorder,” ICD-11 frequently asked questions. WHO defines gaming disorder as a pattern of gaming behavior involving impaired control, increasing priority over other activities, and continuation or escalation despite negative consequences.

⁹⁷ Qur’an 2:219; Qur’an 5:90-91.

⁹⁸ OpenAI, “Expanding on What We Missed with Sycophancy,” May 2, 2025. OpenAI stated that an April 2025 GPT-4o update made the model more sycophantic, including by validating doubts, fueling anger, urging impulsive actions, or reinforcing negative emotions, and noted safety concerns related to mental health, emotional over-reliance, and risky behavior.

that they were right and reduce willingness to repair interpersonal conflict.⁹⁹ This is not cited as religious authority, but as technical evidence that AI systems can be tuned in ways that affect user judgment and emotional regulation. From a fiqh perspective, such systems intensify the need to examine not only content, but also persuasive design and foreseeable psychological effects.

7.6 Negligence Through Overreliance

The sixth category is negligence through overreliance. A Muslim may use AI to assist with low-stakes tasks, but he may not outsource duties that require his own verification, judgment, expertise, or moral responsibility. This is especially important in religious, medical, legal, educational, financial, and family contexts. If a person sends AI-generated advice, claims, citations, or instructions to others without verification, he may become responsible for harm caused by negligent reliance. The Qur'an warns against speaking without knowledge and against pursuing that of which one has no knowledge.¹⁰⁰ This principle applies with special force to AI outputs that are plausible but unverified. Fluency is not knowledge, and confidence is not *bayyinah*.

7.7 The Spread of Harmful AI-Generated Content

The seventh category is the spread of harmful AI-generated content. A person who uses AI improperly may expose others to falsehood, indecency, confusion, or sin. This includes generating sexual material, sharing fabricated claims, spreading extremist propaganda, creating misleading religious content, or circulating AI-generated rumors. The moral responsibility is not limited to the original prompt. In Islam, publication and dissemination matter. One may sin by producing falsehood, by sharing it, by amplifying it, by monetizing it, or by knowingly providing the means for others to act upon it. The Prophet, peace be upon him, warned that whoever calls to misguidance bears a burden like the burdens of those who follow him, without decreasing their burdens at all.¹⁰¹ This principle is highly relevant to scalable AI-generated harm.

7.8 Individual Responsibility and Platform Responsibility

The eighth category is the tension between individual responsibility and platform responsibility. A user cannot excuse deliberate sin by blaming an algorithm. At the same time, companies cannot excuse foreseeable harm by saying users made choices in an environment deliberately engineered to shape those choices. Islamic law recognizes both direct causation and indirect causation, both *mubāsharah* and *tasabbub*, and it recognizes liability where negligence or wrongful facilitation produces harm. The precise distribution of responsibility depends on knowledge, intention, foreseeability, control, safeguards, and causal contribution. A user who seeks prohibited content is sinful for seeking it. A platform that knowingly recommends more of it, profits from it, and fails to implement reasonable safeguards may share responsibility according to its role.

This layered responsibility is especially important when AI systems affect minors and

⁹⁹ Myra Cheng et al., "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence," *Science* 391, no. 6792 (March 26, 2026), <https://doi.org/10.1126/science.aec8352>.

¹⁰⁰ Qur'an 17:36.

¹⁰¹ *Ṣaḥīḥ Muslim*, no. 2674.

vulnerable users. Children, adolescents, grieving people, lonely users, compulsive users, and those experiencing mental health crises may be less able to resist manipulative design or evaluate harmful output. UNESCO's education guidance emphasizes age-appropriate use, human agency, and the protection of data privacy in generative AI education.¹⁰² These principles correspond to Islamic concerns for guardianship, protection from harm, and differential vulnerability. The Shari'ah does not treat a child, an expert adult, an emotionally distressed person, and a trained professional as identical users. Household rules, school policies, and institutional safeguards should therefore be stricter where vulnerability is greater.

7.9 The Principle of *Taklif*

The principle of *taklif* must also be handled carefully. Engineered compulsion, addiction-like design, emotional manipulation, and algorithmic pressure may reduce self-control or increase difficulty, but they do not automatically erase accountability. Allah does not burden a soul beyond its capacity.¹⁰³ Capacity, coercion, forgetfulness, ignorance, and compulsion are all relevant in Islamic law, but ordinary temptation is not the same as legal compulsion. The more a person knowingly enters a harmful environment, ignores warnings, conceals misconduct, and continues despite clear harm, the stronger his personal accountability. The more a system exploits vulnerability, conceals risks, targets minors, or removes meaningful choice, the stronger the responsibility of designers and deployers.

7.10 Practical fiqh framework

A practical ruling framework may therefore be summarized as follows:

1. AI assistance in daily conduct is permissible when it supports a lawful purpose, remains within the expectations of the relevant relationship or institution, does not misrepresent competence or authorship, does not violate confidentiality or rights, and does not lead to prohibited content or harmful dependency.
2. AI use becomes prohibited when it involves cheating, plagiarism, forged credentials, fabricated sources, false reviews, impersonation, fraudulent communication, religious misrepresentation, or concealment of material facts.
3. AI use becomes blameworthy or prohibited when it predictably cultivates compulsive engagement, neglect of obligations, exposure to haram, emotional overreliance, or manipulation of others.
4. Platform and developer responsibility increases when harmful outcomes are foreseeable, scalable, profitable, and insufficiently mitigated.

This section should also guide institutions. Muslim schools and seminaries should define AI use in assessment, distinguishing permitted assistance from prohibited substitution. Employers should clarify whether AI use is allowed, what must be disclosed, and where human review is required. Islamic organizations should avoid publishing AI-generated religious content without qualified review. Parents should set rules for children's AI use, especially around chatbots, search, media generation, and private disclosure. Imams should teach not only "do not cheat," but also the deeper principle: a Muslim's public presentation must correspond to reality where others rely upon it.

The daily fiqh of AI is ultimately a fiqh of amanah. A person's words, work, credentials, promises, relationships, and attention are trusts. AI can help fulfill those trusts when used transparently and responsibly. It can also help betray them with unprecedented

¹⁰² Miao and Holmes (2023), *Guidance for Generative AI in Education and Research*.

¹⁰³ Qur'an 2:286.

ease. The Sharī'ah does not condemn beneficial tools, but it does condemn falsehood, cheating, betrayal, negligence, and cooperation in sin. The Muslim response must therefore be neither panic nor permissiveness, but disciplined moral agency: using assistance without surrendering responsibility, seeking benefit without normalizing deception, and remembering that technological ease does not lighten the weight of accountability before Allah.

8. Practical Guidelines for Daily-Use AI

The preceding analysis shows that daily-use AI cannot be governed by a single ruling that applies uniformly to every user, tool, and context. The more appropriate approach is a layered framework of guidance: personal discipline for individuals, protective structure for families, verification and pastoral care for imams, academic integrity policies for schools, governance protocols for institutions, and moral accountability for developers and platforms. These guidelines do not replace case-specific *iftā'*, but they provide a practical application of the principles already established: *amanah*, *tabayyun*, protection from harm, preservation of family and intellect, avoidance of deception, and consideration of consequences.

8.1 Guidelines for Individuals

For individuals, the first obligation is to preserve moral agency. A Muslim should not use AI in a way that transfers personal responsibility to a system that has no *taqwā*, no accountability before Allah, and no duty to tell the truth. AI may assist lawful activity, but the user remains responsible for what he asks, what he accepts, what he shares, what he conceals, and what he presents as his own. A useful personal rule is that AI may assist effort but should not falsely replace effort where effort is being evaluated. It may improve wording, but not fabricate sincerity. It may help organize thought, but not impersonate knowledge. It may help locate information, but not remove the duty of verification when religion, reputation, money, safety, or rights are involved.

Individuals should therefore distinguish between assistance, reliance, and substitution. Assistance remains permissible in ordinary cases when it serves a lawful purpose and does not mislead others. Reliance becomes more serious when the user acts on AI output without verifying it in matters that affect other people. Substitution becomes impermissible when the user presents AI-generated work, emotion, judgment, authorship, or competence as though it were personally produced in a context where that distinction matters. The Qur'anic command not to pursue what one does not know applies with particular force to an age in which confident output can be produced without personally possessing that knowledge.¹⁰⁴

Individuals should also adopt a rule of disclosure where trust depends on authorship or personal competence. In casual use, disclosure may not matter. In academic, professional, religious, legal, medical, financial, or public-facing settings, disclosure may be required by *amanah* even when no one asks. If the audience, teacher, employer, client, spouse, congregation, or reader reasonably assumes that the work reflects the person's own effort or expertise, undisclosed AI substitution may become deception. This does not mean every grammar correction must be announced, or that AI cannot be used for improving one's writing, projects, or proposals. It means that the higher the reliance placed upon personal authorship, skill, sincerity, or judgment, the stronger the duty of disclosure. In some workplaces today, it is a given that employees are trained on AI and in fact expected to utilize it to improve their final ideas or work.

Individuals should treat AI companionship with disciplined boundaries. A non-erotic tool used for reminders, conversation practice, organization, or limited emotional expression may be permissible. But one should avoid private patterns of attachment that cultivate romantic fantasy, sexual desire, secrecy, emotional dependence, or withdrawal from family and community. A simple personal test is whether one would feel shame if the nature of the

¹⁰⁴ Qur'an 17:36.

interaction were known to a spouse, parent, mentor, or trusted scholar. Shame is not always proof of prohibition, but it can reveal that the interaction has crossed from harmless use into a morally compromised space. The Prophet, peace be upon him, described modesty as part of faith, and the Qur'an commands believers not to approach *zinā*.¹⁰⁵

Individuals must also protect privacy. They should not enter confidential information into AI systems without understanding the risks, especially information involving spouses, children, students, patients, clients, employees, donors, legal matters, medical records, private sins, or religious questions involving identifiable people. Even when a person is willing to disclose his own information, he may not have the right to disclose another's private affairs. This follows from the prohibition of *tajassus*, the protection of dignity, and the duty to preserve trusts.¹⁰⁶

8.2 Guidelines for Families

For families, the central responsibility is not merely technical restriction but moral formation. Parents should not treat AI as harmless entertainment simply because it appears conversational or helpful. Children and adolescents are still developing judgment, impulse control, modesty, social identity, and academic discipline. AI tools that simulate friendship, romance, constant validation, or adult conversation may shape expectations before the child has the maturity to evaluate them. Parents should therefore *fully* know which AI systems their children use, whether those systems include companion features, image generation, private chat, sexual content, voice interaction, or memory, and whether the child is using AI for schoolwork in ways that violate academic trust.

Family guidelines should include age-appropriate rules. Younger children should not have unsupervised access to open-ended AI chatbots or companion systems. Adolescents should be taught that AI can produce false information, manipulate emotions, expose them to inappropriate material, and make cheating easy. Parents should explain not only what is forbidden but why: because the Muslim is trained in honesty, modesty, effort, and accountability. The household should also establish "no AI zones" where human relationships are preserved: meals, family conversations, worship, bedtime routines, and conflict resolution. A child who learns to outsource boredom, study, apology, and companionship to AI may not develop the human virtues those moments were meant to cultivate.

Families should also prepare for difficult cases. If a child or spouse develops attachment to an AI companion, the first response should not be mockery. The attachment may indicate loneliness, anxiety, marital strain, social isolation, pornography exposure, depression, or emotional unmet need. The religious response should combine boundaries with serious care. The family may need to restrict the tool, but it should also address the underlying vulnerability. The goal is not only to remove access but to restore healthy relationships, meaningful activity, worship, counsel, and human support.

8.3 Guidelines for Religious Leadership

For imams and religious counselors, the first guideline is verification. Imams will increasingly encounter screenshots, recordings, AI-generated messages, synthetic images, alleged confessions, fabricated religious clips, and claims about what a person "said" or "did." They should not act on digital evidence without appropriate *tabayyun*, especially

¹⁰⁵ Qur'an 17:32; *ṣaḥīḥ al-Bukhārī*, no. 9.

¹⁰⁶ Qur'an 49:12; Qur'an 17:70; Qur'an 4:58.

where marriage, reputation, employment, criminal allegation, community office, or religious standing is at stake. The more serious the accusation, the stronger the need for corroboration, provenance, expert review, and procedural restraint. The Shari'ah's protection of reputation requires caution before accepting or amplifying claims that may destroy a person's honor.¹⁰⁷

Imams should also teach congregants that forwarding unverified AI-generated content can itself be *potentially* sinful depending on the content. The Prophet, peace be upon him, warned that it is enough falsehood for a person to narrate everything he hears.¹⁰⁸ In an AI environment, "everything he hears" includes everything he sees, receives, screenshots, downloads, or encounters online. Community members should be warned against spreading synthetic scandals, fake scholar statements, fabricated fatwas, manipulated political clips, and AI-generated religious stories without verification.

Religious leaders should not use AI to issue fatwas, answer individualized religious questions, or produce religious content without qualified review. AI may help organize research, locate references, summarize non-authoritative material, or draft administrative communication, but the imam remains responsible for accuracy, even if the platform is an Islamic AI service. Sermons, fatwas, family advice, and public religious guidance require human knowledge, context, accountability, and fear of Allah. If AI is used administratively or editorially, it should remain subordinate to qualified review. The danger is not only factual error, but a gradual erosion of scholarly *amanah* when religious guidance begins to sound authoritative while being generated without understanding.

Imams should also develop pastoral guidance for AI companionship and virtual intimacy. They should avoid sensational language that equates every AI interaction with *zinā*, but they should speak clearly about romantic AI partners, sexualized chatbots, secrecy, marital harm, and dependency. Many congregants may be embarrassed to ask about these matters. A wise imam will teach the principles with wisdom and cautiousness before a crisis occurs: preserving the heart, guarding desire, protecting marriage, avoiding secrecy, and seeking human counsel rather than synthetic validation. The hadith that religion is sincere counsel should guide the pastoral tone.¹⁰⁹

8.4 Guidelines for Educational Institutions

For schools, seminaries, and educational programs, the most urgent need is clear policy. Ambiguity invites both genuine confusion and dishonest exploitation. Institutions should define acceptable, restricted, and prohibited AI use across assignments. For example, AI may be allowed for brainstorming, grammar assistance, translation support, accessibility, or practice questions if disclosed and if the learning objective allows it. It should be prohibited where the assessment is designed to measure the student's own reasoning, memorization, writing, source analysis, Arabic ability, legal reasoning, or research process unless the instructor explicitly permits it. The key question is not simply "Was AI used?" but "Did the use undermine what the assessment was meant to measure?"

Schools should redesign assessment rather than merely police students. Oral exams, in-class writing, annotated bibliographies, process-based drafts, source defense, live problem-solving, personalized reflection, verbal presentations, and supervised research can reduce dishonest use while preserving the benefits of AI literacy. Islamic seminaries in

¹⁰⁷ Qur'an 49:6; Qur'an 24:4; *Ṣaḥīḥ Muslim*, no. 1711a.

¹⁰⁸ *Ṣaḥīḥ Muslim*, no. 5.

¹⁰⁹ *Ṣaḥīḥ Muslim*, no. 55.

particular should teach students when AI can assist research and when it becomes dangerous: fabricated citations, mistranslated sacred texts, decontextualized fatwas, weak hadith claims, false attributions, and confident but inaccurate summaries of madhab positions. UNESCO's guidance on generative AI in education emphasizes institutional governance, privacy, age-appropriate use, and preservation of human agency; Muslim schools should add to this a religious concern for amanah, adab al-talab, and truthful certification.¹¹⁰

Educational institutions should also clarify consequences. If AI cheating is discovered, the response should distinguish between ignorance, negligence, deliberate fraud, repeated misconduct, and high-stakes certification. A first-year student who misunderstood a policy is not identical to a graduate student fabricating sources in a thesis or a seminary student submitting AI-generated legal analysis as his own. Disciplinary policy should be firm but proportionate, combining correction, education, restitution where needed, and protection of institutional integrity.

8.5 Guidelines for Islamic Institutions

For Muslim institutions more broadly, including Islamic centers, charities, publishers, counseling centers, Islamic schools, and seminaries, AI governance should become a formal matter of policy. Institutions should identify which uses are allowed, who may approve them, what data may be entered, when human review is required, and what kinds of AI-generated material may be published. They should prohibit entering confidential counseling notes, donor data, student records, marital disputes, legal information, private sins, or identifiable religious questions into public AI tools unless there is a secure, approved, and legally compliant system with proper consent and oversight.

Institutions should also adopt a synthetic media policy. Official communications should not use a scholar's voice, image, likeness, or writing style synthetically without explicit consent and clear disclosure. Charities should not generate fake beneficiary stories or images. Masajid should not publish AI-generated religious videos that could be mistaken for real scholarly speech. Schools should not use AI images of students without parental permission. Publishers should disclose AI-generated or AI-assisted material where authorship, originality, or trust is relevant. In religious contexts, trust is not a decorative value; it is the foundation of authority.

Institutional governance should consider including risk tiers. Low-risk uses such as grammar assistance, scheduling, formatting, and non-confidential administrative drafting may require minimal oversight. Medium-risk uses such as public communication, educational content, student support, or donor communication require human review and disclosure norms. High-risk uses such as religious guidance, counseling, disciplinary decisions, hiring, expulsion, legal claims, deepfake detection, or crisis response require qualified human authority and should not be delegated to AI. This tiered approach corresponds to the Islamic principle that rulings and safeguards vary with the severity and probability of harm.

Institutions should also create a response protocol for AI-related harm. This includes procedures for handling deepfake allegations, synthetic sexual abuse, fake religious clips, AI-assisted plagiarism, impersonation, and data exposure. A serious protocol should include preservation of evidence, avoidance of public accusation before verification, consultation

¹¹⁰ Fengchun Miao and Wayne Holmes, *Guidance for Generative AI in Education and Research* (Paris: UNESCO, 2023).

with technical experts where needed, pastoral support for victims, reporting obligations where applicable, and public correction when false content has spread. NIST's AI Risk Management Framework is useful here as a secular risk-management reference because it emphasizes mapping, measuring, managing, and governing AI risk; Islamic institutions can adapt this structure while grounding their objectives in Sharī'ah values.¹¹¹

8.6 Guidelines for Developers and Technologists

For developers and technologists, the guidance must go beyond personal piety. Building systems that affect millions of users is a form of public trust. The Qur'an commands cooperation in righteousness and forbids cooperation in sin and transgression.¹¹² This applies not only to users who request harmful content but also to those who design systems that predictably facilitate deception, sexual exploitation, addiction-like engagement, privacy violation, or reputational harm. A Muslim developer should ask not only whether his application or contribution to an application is lawful in isolation, but what use case it enables, what harm it scales, what safeguards it lacks, and whether he is contributing to a system whose business model depends on human weakness.

Developers should treat foreseeability as morally relevant. If a model can easily generate nonconsensual sexual imagery, impersonate real people, fabricate religious speech, groom minors, encourage self-harm, produce fraudulent documents, or manipulate emotionally vulnerable users, then safeguards are not optional. They are part of amanah. The Prophetic warning regarding one who practices medicine without recognized competence and causes harm is often cited by jurists as a basis for professional liability.¹¹³ The analogy is not exact, but the principle is clear: specialized skill that can harm others carries heightened responsibility. AI developers and deployers cannot hide behind the claim that "users decide" when harmful use is predictable, scalable, and insufficiently mitigated.

Developers should also avoid designing for dependency. Engagement metrics are not morally neutral when they reward compulsion, outrage, sexual escalation, emotional dependence, or endless distraction. A system that manipulates vulnerabilities in order to maximize retention may be profitable while remaining ethically corrupt. Muslim technologists should advocate for friction, transparency, age restrictions, crisis escalation, content labeling, user controls, privacy protection, auditability, and meaningful consent. These are not merely secular compliance concerns; they are applications of preventing harm, preserving dignity, and refusing cooperation in sin. Furthermore, Muslim technologists should be aware that sometimes age restrictions and transparency are Trojan horses for normalizing surveillance and repressing speech uncomfortable to entrenched money and power such as whistleblowing and dissent;¹¹⁴ advocacy should focus on the end outcome (removing dependency as an application objective) in such a way so as to not

¹¹¹ National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (Gaithersburg, MD: NIST, January 2023).

¹¹² Qur'an 5:2.

¹¹³ *Sunan Abī Dāwūd*, no. 4586; *Sunan al-Nasā'ī*, no. 4830; *Sunan Ibn Mājah*, no. 3466. The hadith is transmitted through multiple routes, classified as ḥasan by Al-Albani, and is cited by jurists in discussions of liability for unqualified medical practice.

¹¹⁴ Eric Goldman, "The 'Segregate-and-Suppress' Approach to Regulating Child Safety Online," *Stanford Technology Law Review* 28 (2025): 173; David E. Pozen, "Transparency's Ideological Drift," *Yale Law Journal* 128, no. 1 (2018): 100; Todd W. Shaw, "When Text and Policy Conflict: Internal Whistleblowing under the Shadow of Dodd-Frank," *Administrative Law Review* 70, no. 3 (2018): 673–96.

inadvertently facilitate other types of harms or friction and empower corrupt power structures.

Developers working inside large companies may not control every product decision. Their responsibility varies with role, knowledge, authority, proximity to harm, and ability to influence outcomes. A junior engineer maintaining neutral infrastructure is not identical to a product lead designing sexualized companion bots or a growth team optimizing harmful content recommendations. But limited authority does not remove all responsibility. A Muslim professional should document concerns, advocate internally, seek safer assignments where possible, refuse direct work on clearly prohibited applications, and leave work that requires direct participation in grave harm when no lawful alternative within the role remains. This requires case-specific fiqh, because employment, necessity, and degree of involvement differ.

Across all categories, the most important practical principle is that AI should remain subordinate to accountable human judgment. It may assist, but it should not become a concealed author, fabricated witness, artificial spouse, unqualified mufti, manipulative counselor, automated accuser, or substitute conscience. The Sharī'ah does not demand that Muslims abandon every beneficial technology. It demands that benefit remain governed by truth, dignity, modesty, justice, and accountability. These guidelines are therefore not merely defensive. They are an attempt to preserve the conditions under which Muslims can use powerful tools without allowing those tools to deform the moral life those tools were meant to serve.

This paper has discussed *amanah* as it relates to individuals using AI tools in the context of epistemic discipline, privacy, moral environment, deepfakes, transmission of religion/isnad, general public amanah, deception of competence and other professional deception. Each of these areas is in need of tools developed with the moral objectives outlined in this paper in contraposition to most modern commercial objectives of engagement optimization and dependency maximization. Tools that preserve epistemic integrity, make it easy to accomplish beneficial aims and difficult to achieve *amanah*-hurting ones will benefit not just Muslims but humanity at large. It behooves Muslims today to take the lead in developing these technologies.

9. Conclusion

Daily-use artificial intelligence is not a single moral problem. It is a family of technologies that now mediate how people speak, study, desire, remember, trust, verify, work, worship, and relate to one another. Its ethics cannot be viewed simplistically by celebration nor suspicion. It may assist lawful learning, accessibility, productivity, creativity, translation, organization, and communication. It may also facilitate deception, unlawful intimacy, reputational injury, privacy violation, dependency, false attribution, academic dishonesty, and the weakening of moral agency. A sound Sunni fiqh approach must therefore avoid two opposite errors: treating AI as inherently corrupt because it is new, and treating it as morally neutral because it is technically a tool.

The Sharī'ah's approach to human action is neither technologically naïve nor reactionary. It begins from the accountability of the mukallaf, the protection of rights, the preservation of truth, the fulfillment of trusts, the avoidance of harm, and the pursuit of benefit within divine limits. AI has no *taqwā*, no *dhimmah*, no *niyyah*, no *wilayah*, and no accountability before Allah. The machine is not the moral subject. Yet the human beings who design, deploy, market, use, rely upon, conceal, share, and profit from AI remain accountable for their actions and omissions. This distinction is foundational. It prevents the anthropomorphizing of AI while also preventing humans from hiding behind the machine.

A proactive Muslim approach to AI must therefore be principled rather than merely defensive. The community should not wait until habits harden, children are shaped by synthetic companions, reputations are destroyed by fabricated media, students normalize deception, and religious speech is polluted by artificial attribution; nor should it reject beneficial tools that serve education, accessibility, research, communication, and social welfare. The proper path is earlier ethical formation, clearer policies, better technical literacy, stronger verification norms, and a renewed culture of *amānah*.

This is also why AI ethics cannot be left entirely to market incentives or secular risk frameworks: such models may help identify privacy risks, transparency needs, and system-level harms, but they cannot supply the moral vocabulary of worship, chastity, *taqwā*, *amānah*, *ḥayā'*, *barakah*, sin, repentance, and accountability before Allah. The Sharī'ah does not ask Muslims to be technologically primitive; it asks them to be morally awake. The Muslim contribution to AI ethics should therefore extend beyond asking whether Muslims may use tools created by others, to asking what kinds of tools should exist, what human goods they should serve, and what harms they must never normalize.

The conclusion, then, is one of disciplined caution and constructive engagement. Daily-use AI should be welcomed where it assists lawful benefit, expands access, reduces hardship, and remains subordinate to accountable human judgment. It should be restricted where it weakens moral agency, violates privacy, corrupts intimacy, deceives audiences, fabricates evidence, damages reputation, or exploits vulnerability. It should be prohibited where it becomes a direct means to falsehood, sexualized exploitation, impersonation, fraud, religious misguidance, or serious harm. These distinctions allow Islamic law to be both firm and fair, protective and practical, rooted and responsive.

The coming years will require more detailed research on liability, platform governance, minors, mental health, digital evidence, biometric rights, AI in religious education, and the professional obligations of Muslim technologists. The rulings offered here should therefore be treated as a framework for ongoing *iftā'*, not the final word on every emerging use case. As AI systems become more personalized, multimodal, autonomous, and embedded in daily life, scholars will need sustained collaboration with technologists,

educators, clinicians, legal experts, and community leaders. Such collaboration is not a dilution of fiqh. It is part of *fiqh al-wāqī'* and the responsible fulfillment of scholarly amanah.

Islamic law is capable of meeting this moment without imitation or panic. Its sources command truthfulness, protect dignity, regulate desire, preserve family, require verification, condemn deception, and hold human beings accountable for the harms they cause or facilitate. Its *uṣūlī* methods allow scholars to classify new realities, weigh competing interests, and issue rulings proportionate to context. Its spiritual vision reminds believers that the most important question is not whether a machine can imitate intelligence, but whether human beings will preserve wisdom, restraint, and accountability while using it. In an age of artificial fluency, synthetic intimacy, and manufactured certainty, the Muslim task is to remain truthful, dignified, discerning, and answerable to Allah.

Limitations of This Paper

The scope of this paper is deliberately limited to daily-use AI as it affects ordinary Muslim conduct, social relations, privacy, truthfulness, emotional attachment, self-presentation, and moral harm. It does not attempt to resolve all questions in the fiqh of artificial intelligence. Autonomous weapons, state surveillance, predictive policing, judicial AI, medical triage, autonomous vehicles, Islamic finance applications, employment in AI companies, dual-use engineering, AI-assisted *iftā'*, AI-generated fatwas, and the institutional governance of high-stakes AI each require separate treatment because they involve different levels of authority, expertise, risk, liability, and public consequence.

These limits are methodological, not evasive. A daily-use AI tool may become high-stakes when it is used to generate evidence, imitate a scholar, advise a vulnerable person, influence marital conduct, assist formal certification, produce religious claims, expose private information, or create reputational harm. In such cases, the legal analysis must rise to the level of the rights and consequences involved. The purpose of this paper is therefore to provide a framework for ordinary and emerging cases, not a final ruling for every possible AI deployment. It offers principles, classifications, and practical guidance for *iftā'* while recognizing that new applications will require continued collaboration among jurists, technologists, educators, psychologists, legal experts, and community leaders. This limitation is itself part of scholarly amanah: one must neither ignore urgent harms because the field is new, nor issue overbroad rulings beyond the case actually studied.

Bibliography

- American Psychological Association. "Use of Generative AI Chatbots and Wellness Applications for Mental Health." APA Health Advisory, 2025.
- Associated Press. "Trump Highlights Partnership Investing \$500 Billion in AI." January 21, 2025.
- Abū Dāwūd, Sulaymān ibn al-Ash'ath al-Sijistānī. *Sunan Abī Dāwūd*.
- Auda, Jasser. *Maqasid al-Shariah as Philosophy of Islamic Law: A Systems Approach*. London: International Institute of Islamic Thought, 2008.
- Balalle, Himendra, and Sachini Pannilage. "Reassessing Academic Integrity in the Age of AI: A Systematic Literature Review on AI and Academic Integrity." *Social Sciences & Humanities Open* 11 (2025): 101299. <https://doi.org/10.1016/j.ssaho.2025.101299>.
- Brandtzaeg, Petter Bae, Marita Skjuve, and Asbjørn Følstad. "My AI Friend: How Users of a Social Chatbot Understand Their Human-AI Friendship." *Human Communication Research* 48, no. 3 (2022): 404–29.
- Bukhārī, Muḥammad ibn Ismā'īl al-. *Ṣaḥīḥ al-Bukhārī*.
- Campbell, Laurie O., Kathryn Babb, Glenn W. Lambie, and B. Grant Hayes. "An Examination of Generative AI Response to Suicide Inquires [sic]: Content Analysis." *JMIR Mental Health* 12 (2025): e73623. <https://doi.org/10.2196/73623>.
- Cheng, Myra, et al. "Sycophantic AI Decreases Prosocial Intentions and Promotes Dependence." *Science* 391, no. 6792 (March 26, 2026). <https://doi.org/10.1126/science.aec8352>.
- Chesney, Robert, and Danielle Keats Citron. "Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security." *California Law Review* 107, no. 6 (2019): 1753–1820.
- Cleveland-Stout, Nick. "Israel Wants to Train ChatGPT to Be More Pro-Israel." *Responsible Statecraft*, September 29, 2025.
- Common Sense Media. *Social AI Companions: AI Risk Assessment*. San Francisco: Common Sense Media, 2025.
- Davenport, Thomas H., and John C. Beck. *The Attention Economy: Understanding the New Currency of Business*. Boston: Harvard Business School Press, 2001.
- De Freitas, Julian, Zeliha Oğuz-Uğuralp, Ahmet Kaan Uğuralp, and Stefano Puntoni. "AI Companions Reduce Loneliness." *Journal of Consumer Research* 52, no. 6 (2026): 1126–48. <https://doi.org/10.1093/jcr/ucaf040>.
- De Freitas, Julian, Zeliha Oğuz-Uğuralp, and Ahmet Kaan-Uğuralp. "Emotional Manipulation by AI Companions." arXiv preprint arXiv:2508.19258, August 15, 2025, revised October 2025.
- Edelson, Daniel, and Raphael Kahan. "Report: Israel to Spend Over Half a Billion Shekels Turning ChatGPT into Public Diplomacy Tool." *Ynetnews*, October 6, 2025.
- European Union. Regulation (EU) 2024/1689, Artificial Intelligence Act. 2024.
- Fitzpatrick, Kathleen Kara, Alison Darcy, and Molly Vierhile. "Delivering Cognitive Behavior Therapy to Young Adults with Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent: A Randomized Controlled Trial." *JMIR Mental Health* 4, no. 2 (2017): e19.
- Fogg, B. J. *Persuasive Technology: Using Computers to Change What We Think and*

- Do. San Francisco: Morgan Kaufmann, 2003.
- Ghazālī, Abū Ḥāmid al-. *Iḥyā' 'Ulūm al-Dīn*. Vol. 3. Beirut: Dār al-Ma'rifah, n.d.
- Goldman, Eric. "The 'Segregate-and-Suppress' Approach to Regulating Child Safety Online." *Stanford Technology Law Review* 28 (2025): 173.
- Goldstein, Josh A., and Andrew Lohn. "Deepfakes, Elections, and Shrinking the Liar's Dividend." Brennan Center for Justice, January 23, 2024.
- Hani, Suleiman. *Fiqh al-Muwāzanat: The Jurisprudence of Weighing Benefits and Harms*. Waymark Publications, 2026.
- He, Yuhao, et al. "Conversational Agent Interventions for Mental Health Problems: Systematic Review and Meta-analysis of Randomized Controlled Trials." *Journal of Medical Internet Research* 25 (2023): e43862.
- Ho, Joyce Q. H., et al. "Potential and Pitfalls of Romantic Artificial Intelligence (AI) Companions: A Systematic Review." *Computers in Human Behavior: Artificial Humans* 3 (2025): 100130.
- Hu, Krystal, and Stephen Nellis. "Trump Announces Private-Sector \$500 Billion Investment in AI Infrastructure." *Reuters*, January 21, 2025.
- Ḥusayn, Walīd b. 'Alī al-. *I'tibār Ma'ālāt al-Af'āl wa-Atharuhā al-Fiqhī*. 2 vols. Riyadh: Dār al-Tadmuriyyah, 2009.
- Ibn 'Ābidīn, Muḥammad Amīn. *Radd al-Muḥtār*. Beirut: Dār al-Fikr, 1992.
- Ibn al-Qayyim al-Jawziyyah. *I'lām al-Muwaqqi'īn 'an Rabb al-'Ālamīn*. Edited by Muḥammad 'Abd al-Salām Ibrāhīm. 4 vols. Beirut: Dār al-Kutub al-'Ilmiyyah, 1991.
- Ibn 'Āshūr, Muḥammad al-Tāhir. *Maqāṣid al-Sharī'ah al-Islāmiyyah*. Edited by El-Tahir el-Mesnawi. Kuala Lumpur: Al-Fajr, 1999.
- Ibn Mājah, Muḥammad ibn Yazīd. *Sunan Ibn Mājah*.
- Ibn Miskawayh, Abū 'Alī Aḥmad ibn Muḥammad. *Tahdhīb al-Akhlāq wa-Taṭhīr al-A'rāq*. Edited by Ibn al-Khaṭīb. Cairo: al-Maṭba'ah al-Miṣriyyah wa-Maktabatuhā, n.d.
- Ibn Nujaym. *Al-Ashbāh wa-l-Nazā'ir*. Beirut: Dār al-Kutub al-'Ilmiyyah, n.d.
- Ibn Rushd. *Bidāyat al-Mujtahid wa-Nihāyat al-Muqtaṣid*. Al-Bābī al-Ḥalabī edition.
- Ibn Taymiyyah. *Dar' Ta'āruḍ al-'Aql wa-l-Naql*. Edited by Muḥammad Rashād Sālim. 11 vols. Riyadh: Jāmi'at al-Imām Muḥammad b. Sa'ūd al-Islāmiyyah, 1411/1991.
- Ibn Taymiyyah. *Majmū' al-Fatāwā*. Edited by 'Abd al-Raḥmān ibn Qāsim. 37 vols. Al-Manṣūrah: Dār al-Wafā', 2005.
- International Islamic Fiqh Academy. Resolution No. 43 (5/5) on Moral Rights.
- International Islamic Fiqh Academy. Resolution No. 258 (3/26), "Artificial Intelligence: Its Rulings, Guidelines, and Ethics."
- Jadīyah, 'Umar. *Aṣl I'tibār al-Ma'āl bayna al-Nazariyyah wa-l-Taṭbīq*. Beirut: Dār Ibn Ḥazm, 2009.
- Jones, Bethany A., and Camille Nurka. "Labiaplasty and Pornography: A Preliminary Investigation." *Porn Studies* 2, no. 1 (2015): 62–75.
- Jordanian General Iftaa' Department. "Ruling on Converting Photographs into Drawings Using Artificial Intelligence Technology." Fatwa no. 3989, July 24, 2025.
- Kamālī, 'Abd Allāh Yaḥyā al-. *Ta'sīl Fiqh al-Muwāzanāt*. Beirut: Dār Ibn Ḥazm, 2000.
- Kasneci, Enkelejda, et al. "ChatGPT for Good? On Opportunities and Challenges of

- Large Language Models for Education." *Learning and Individual Differences* 103 (2023): 102274.
- Kira, Beatriz. "When Non-Consensual Intimate Deepfakes Go Viral." *Computer Law & Security Review* 54 (2024): 106024.
- Laestadius, Linnea, et al. "Too Human and Not Human Enough: A Grounded Theory Analysis of Mental Health Harms from Emotional Dependence on the Social Chatbot Replika." *New Media & Society* 26, no. 10 (2024): 5923–41.
- Li, Han, et al. "Systematic Review and Meta-Analysis of AI-Based Conversational Agents for Promoting Mental Health and Well-Being." *npj Digital Medicine* 6 (2023): 236.
- Liang, Weixin, et al. "GPT Detectors Are Biased Against Non-Native English Writers." *Patterns* 4, no. 7 (2023): 100779.
<https://doi.org/10.1016/j.patter.2023.100779>.
- Lorenz, Taylor. "An Influencer's AI Clone Will Be Your Girlfriend for \$1 a Minute." *Washington Post*, May 13, 2023.
- Mālik ibn Anas. *Al-Muwatta`*.
- Maples, Bethanie, Merve Cerit, Aditya Vishwanath, and Roy Pea. "Loneliness and Suicide Mitigation for Students Using GPT3-Enabled Chatbots." *npj Mental Health Research* 3 (2024): article no. 4.
- Mardāwī, al-. *Al-Taḥbīr Sharḥ al-Taḥrīr*.
- Miao, Fengchun, and Wayne Holmes. *Guidance for Generative AI in Education and Research*. Paris: UNESCO, 2023.
- Mozilla Foundation. "Creepy.exe: Mozilla Urges Public to Swipe Left on Romantic AI Chatbots Due to Major Privacy Red Flags." February 14, 2024.
- Muslim ibn al-Ḥajjāj. *Ṣaḥīḥ Muslim*.
- Nadwī, `Ali Aḥmed al-. *Al-Qawā'id al-Fiqhiyya*. Damascus: Dār al-Qalam, 1987.
- Nasā'ī, Aḥmad ibn Shu'ayb al-. *Sunan al-Nasā'ī*.
- National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. Gaithersburg, MD: NIST, January 2023.
- National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1. Gaithersburg, MD: NIST, July 2024.
- National Institute of Standards and Technology. *Reducing Risks Posed by Synthetic Content: An Overview of Technical Approaches to Digital Content Transparency*. NIST AI 100-4. Gaithersburg, MD: NIST, 2024.
- Nawawī, Yahyā ibn Sharaf al-. *Al-Arba'ūn al-Nawawīyyah*.
- Nissenbaum, Helen. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford Law Books, 2010.
- OpenAI. "Expanding on What We Missed with Sycophancy." May 2, 2025.
- OpenAI. "New AI Classifier for Indicating AI-Written Text." January 31, 2023. Updated July 2023.
- Organisation for Economic Co-operation and Development. *Recommendation of the Council on Artificial Intelligence*. OECD/LEGAL/0449. Adopted May 22, 2019. Amended May 3, 2024.
- Ortutay, Barbara. "President Trump Signs Take It Down Act, Addressing Nonconsensual Deepfakes. What Is It?" *Associated Press*, May 19, 2025.
- Paslakis, Georgios, et al. "Associations between Pornography Exposure, Body Image

- and Sexual Body Image: A Systematic Review." *Journal of Health Psychology* 27, no. 3 (2022): 743–60.
- Perez, Sarah. "AI Companion Apps on Track to Pull in \$120M in 2025." *TechCrunch*, August 12, 2025.
- Perry, Samuel L., and Cyrus Schleifer. "Till Porn Do Us Part? A Longitudinal Examination of Pornography Use and Divorce." *Journal of Sex Research* 55, no. 3 (2018): 284–296.
- Pozen, David E. "Transparency's Ideological Drift." *Yale Law Journal* 128, no. 1 (2018): 100–165.
- Qarāfī, Shihāb al-Dīn al-. *Al-Furūq / Anwār al-Burūq fī Anwā' al-Furūq*.
- Qarāfī, Shihāb al-Dīn al-. *Sharḥ Tanqīḥ al-Fuṣūl*.
- Qarāḍāwī, Yūsuf al-. *Dirāsāt fī Fiqh Maqāṣid al-Sharī'ah: Bayna al-Maqāṣid al-Kullīyyah wa-l-Nuṣūṣ al-Juz'īyyah*. Cairo: Dār al-Shurūq, 2006.
- Qarāḍāwī, Yūsuf al-. *Fī Fiqh al-Awlawiyyāt: Dirāsah Jadīdah fī Ḍaw' al-Qur'ān wa-l-Sunnah*. Cairo: Maktabat Wahbah, 1996.
- Qurṭubī, Muḥammad ibn Aḥmad al-. *Al-Jāmi' li-Aḥkām al-Qur'ān*.
- Reuters. "Photo from Harris-Walz Rally in Michigan Was Not AI-Generated." August 12, 2024.
- Reuters. "Trump Revokes Biden Executive Order on Addressing AI Risks." January 21, 2025.
- Russell, Stuart, and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 4th ed. Hoboken, NJ: Pearson, 2021.
- Sanūsī, 'Abd al-Raḥmān b. Ma'mar al-. *I'tibār al-Ma'ālāt wa-Murā'āt Natā'ij al-Taṣarrufāt: Dirāsah Muqāranah fī Uṣūl al-Fiqh wa-Maqāṣid al-Sharī'ah*. Dammām: Dār Ibn al-Jawzī, 1424 AH/2003 CE.
- Schiff, Kaylyn Jackson, Daniel S. Schiff, and Natália S. Bueno. "The Liar's Dividend: Can Politicians Claim Misinformation to Evade Accountability?" *American Political Science Review* 119, no. 1 (2025): 71–90.
<https://doi.org/10.1017/S0003055423001454>.
- Sculley, Jordan, and Christopher D. Watkins. "The Great Porn Experiment V2.0: Sexual Arousal Reduces the Salience of Familiar Women When Heterosexual Men Judge Their Attractiveness." *Archives of Sexual Behavior* 51, no. 6 (2022): 3071–82.
- Sharp, Gemma, Marika Tiggemann, and Sarah Mattiske. "Predictors of Consideration of Labiaplasty: An Extension of the Tripartite Influence Model of Beauty Ideals." *Psychology of Women Quarterly* 39, no. 2 (2015): 182–93.
- Shāṭibī, Abū Ishāq al-. *Al-I'tisām*. Beirut: Dār Ibn 'Affān, 1997.
- Shāṭibī, Abū Ishāq al-. *Al-Muwāfaqāt*.
- Shaw, Todd W. "When Text and Policy Conflict: Internal Whistleblowing under the Shadow of Dodd-Frank." *Administrative Law Review* 70, no. 3 (2018): 673–96.
- Shirbīnī, Muḥammad al-Khaṭīb al-. *Mughnī al-Muḥtāj*.
- Shukrī, Farīd. *Fiqh al-Tanzīl*.
- Simon, Herbert A. "Designing Organizations for an Information-Rich World." In *Computers, Communications, and the Public Interest*, edited by Martin Greenberger, 37–72. Baltimore: Johns Hopkins Press, 1971.
- Singh, Kanishka. "OpenAI Wins \$200 Million US Defense Contract." *Reuters*, June 16,

- 2025.
- Skjuve, Marita, Asbjørn Følstad, Knut Inge Fostervold, and Petter Bae Brandtzaeg. "A Longitudinal Study of Human–Chatbot Relationships." *International Journal of Human-Computer Studies* 168 (2022): 102903.
<https://doi.org/10.1016/j.ijhcs.2022.102903>.
- Stanford Institute for Human-Centered Artificial Intelligence. *Artificial Intelligence Index Report 2025*. Stanford, CA: Stanford University, 2025.
- Subkī, Tāj al-Dīn al-. *Jam' al-Jawāmi'*, with al-Bannānī's ḥāshiyah.
- Suyūṭī, Jalāl al-Dīn al-. *Al-Ashbāh wa-l-Nazā'ir fī Qawā'id wa-Furū' Fiqh al-Shāfi'iyyah*. Beirut: Dār al-Kutub al-Ilmiyyah, 1403/1983.
- Talūt, Jamīlah Ḥasan. *Fiqh al-Tanzīl 'inda al-Imām Ibn Taymiyyah*.
- Tong, Anna. "AI Chatbot Company Replika Restores Erotic Roleplay for Some Users." *Reuters*, March 25, 2023.
- Tong, Anna, and Krystal Hu. "What Happens When Your AI Chatbot Stops Loving You Back?" *Reuters*, March 18, 2023.
- UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO, 2021.
- Uriawan, Wisnu, et al. "Preventing AI Deepfake Abuse: An Islamic Ethics Framework." arXiv preprint arXiv:2512.17218, 2025.
<https://doi.org/10.48550/arXiv.2512.17218>.
- Verdoliva, Luisa. "Media Forensics and DeepFakes: An Overview." *IEEE Journal of Selected Topics in Signal Processing* 14, no. 5 (2020): 910–32.
- White House. "Winning the AI Race: America's AI Action Plan." July 2025.
- Whittaker, Joe, Seán Looney, Alastair Reed, and Fabio Votta. "Recommender Systems and the Amplification of Extremist Content." *Internet Policy Review* 10, no. 2 (2021).
- Winner, Langdon. "Do Artifacts Have Politics?" *Daedalus* 109, no. 1 (Winter 1980): 121–36.
- World Health Organization. "Gaming Disorder." ICD-11 frequently asked questions.
- Wright, Paul J., Robert S. Tokunaga, Ashley Kraus, and Elyssa Klann. "Pornography Consumption and Satisfaction: A Meta-Analysis." *Human Communication Research* 43, no. 3 (2017): 315–43.
- Wu, Tim. *The Attention Merchants: The Epic Scramble to Get Inside Our Heads*. New York: Alfred A. Knopf, 2016.
- Yahoo Finance. "The \$1-a-Minute 'Virtual Girlfriend' Bot Has Seen a Surge in Users Since Its Launch and Made \$72,000 in a Week." May 12, 2023.
- Yesilada, M., and S. Lewandowsky. "Systematic Review: YouTube Recommendations and Problematic Content." *Internet Policy Review* 11, no. 1 (2022).
- Zarqā', Muṣṭafā al-. *Al-Madkhal al-Fiqhī al-Āmm*. Damascus: Dār al-Qalam, 2004.
- Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs, 2019.