

THE AI MAQASID FRAMEWORK: INTEGRATING QURANIC OBJECTIVES INTO IR4.0 ERA TEACHING AND GOVERNANCE

Md Muntaqim Meherabⁱ, Md Atiqur Rahmanⁱⁱ, Nuruzzaman Faruquiⁱⁱⁱ, S. M. Aminul Haque^{iv}, Sheak Rashed Haider Noori^v, Tanvirul Islam^{vi}, Belayet Hossain^{vii}, Md Shohidul Islam Polash^{viii}

ⁱ (Corresponding author). B.Sc. (CSE) Student, Department of Computer Science & Engineering, Daffodil International University, Bangladesh. Email: meherab2305101354@diu.edu.bd

ⁱⁱ M.Phil, Al Tafseer Wa Ulumul Quran, Umm Al-Qura University, Makkah, Saudi Arabia.
Email: imamdsc@daffodilvarsity.edu.bd

ⁱⁱⁱ Assistant Professor, Department of Software Engineering, Daffodil International University, Bangladesh.
E-mail: faruqui.swe@diu.edu.bd

^{iv} Professor & Associate Head, Department of Computer Science & Engineering, Daffodil International University, Bangladesh. E-mail: aheadcse2@daffodilvarsity.edu.bd

^v Professor & Head, Department of Computer Science & Engineering, Daffodil International University, Bangladesh. E-mail: headcse@daffodilvarsity.edu.bd

^{vi} Lecturer, Department of Computer Science & Engineering, Daffodil International University, Bangladesh.
E-mail: tanvirulislam.cse@diu.edu.bd

^{vii} Arabic Language, Mahadul Lugah, Al-Azhar University, Cairo, Egypt. E-mail: hfbelayet@yahoo.com

^{viii} Ph.D. Student, Department of Information and Computer Science, King Fahd University of Petroleum and Minerals, Saudi Arabia. E-mail: g202515970@kfupm.edu.sa

ABSTRACT

We present the AI–Maqasid Framework (AIMF), a practical bridge from Qur’anic objectives (al-Maqāsid) to concrete AI design and governance controls for Islamic education in the Fourth Industrial Revolution (IR4.0). AIMF maps the core darūriyyāt—preservation of dīn, nafs, ‘aql, nasl, and māl—to actionable safeguards across the AI lifecycle and cross-walks them with contemporary governance (NIST AI RMF, OECD Principles, UNESCO, EU AI Act). We instantiate AIMF in two education use cases: an Adaptive Tahfiz Tutor and a Qur’anic Studies Assistant based on retrieval-augmented generation (RAG). In both, AIMF yields measurable protections including privacy by design, fairness audits, calibration with uncertainty communication, transparency notes, and Shariah governance checklists. We also provide a brief implementation roadmap, a practical review rubric, and ready-to-use templates to help educators adopt these tools responsibly while safeguarding textual fidelity, learner dignity, and pedagogical authority.

Keywords: Maqasid al-Qur’an, Islamic education, AI governance, IR4.0, retrieval-augmented generation (RAG).

INTRODUCTION

Artificial intelligence has entered classrooms, research labs, and homes, shaping how Muslims learn, teach, and disseminate Qur'anic knowledge[6, 1] The Fourth Industrial Revolution (IR4.0) accelerates this shift through pervasive connectivity and intelligent automation.[7] While global governance (e.g., UNESCO's ethics recommendation, OECD AI Principles, NIST AI RMF, the EU AI Act) codifies safeguards,[8, 9, 10, 11, 2] Islamic education needs a framework that begins from the Qur'an's purposes (Maqāsid)[12, 13, 14, 15] and flows down into technical and institutional controls. This paper addresses that gap.

1. Qur'anic grounding.

In the Holy Qur'an, Allah has said:

O you who believe, be steadfast in upholding justice, bearing witness for God

(Quran 4:135).

O you who believe, if a troublemaker brings you news, investigate

(Qur'an 49:6).

Do not pursue that of which you have no knowledge; surely the hearing, the sight, and the heart each will be questioned

(Qur'an 17:36)

2. Prophetic guidance.

In the Sunnah, the Prophet (peace and blessings be upon him) emphasized truthfulness, mercy, responsibility, and careful speech principles directly relevant to educational AI:

Whoever believes in Allah and the Last Day, let him speak good or remain silent.

(Ṣaḥīḥ al-Bukhārī; Ṣaḥīḥ Muslim)

Whoever deliberately lies about me, let him take his seat in the Fire.

(Ṣaḥīḥ al-Bukhārī; Ṣaḥīḥ Muslim).

He who does not show mercy will not be shown mercy.

(Ṣaḥīḥ al-Bukhārī; Ṣaḥīḥ Muslim).

Each of you is a shepherd, and each of you will be asked about his flock.

(Ṣaḥīḥ al-Bukhārī; Ṣaḥīḥ Muslim).

Guided by these texts, we operationalize:

- Justice (4:135) → parity testing and fairness audits in scoring and retrieval.
- Verification (49:6) → authenticated sources, provenance and citations for every answer.
- Epistemic humility (17:36) → calibration, uncertainty bands, and principled refusals.
- Responsible speech (Bukhārī/Muslim) → harm filters, child-safe defaults, escalation paths.
- Accountability (Bukhārī/Muslim “shepherd” ḥadīth) → Shariah governance board and audit logs.

Why now. IR4.0 has normalized data-driven personalization and ubiquitous automation in classrooms. That scale raises questions that are not merely technical: what to include, what to exclude, and how to preserve sanctity and dignity while using modern tools. AIMF keeps the starting point anchored in Qur'anic aims and then walks down to governance and design controls, so institutions can move fast without drifting from first principles.

Contributions and artefacts.

1. Conceptual: We formalize the AI–Maqasid Framework (AIMF) that maps Qur'anic objectives to concrete AI design and governance controls.
2. Practical: We instantiate AIMF in two education use cases Adaptive Tahfiz Tutor and Qur'anic Studies RAG Assistant (retrieval-augmented, tafsīr-grounded).
3. Operational: We provide metrics, checklists, and a deployment roadmap, and release ready-to-use artefacts:
 - controls catalogue keyed to Maqāsid;
 - lifecycle gates with audit artefacts;
 - risk–control map for both use cases;
 - expert review rubric;
 - practitioner templates (transparency notes; scholar review form).

What is new. Unlike general-purpose AI governance frameworks AIMF starts from the Quran's higher aims and turns them into auditable, classroom-ready controls for Islamic education.

Positioning and novelty. AIMF closes a practical gap between high-level ethics and day-to-day Qur'anic pedagogy. In short:

1. Qur'an-first translation layer: We map al-Maqāsid directly to enforceable controls (authenticated sources, doctrinal boundaries, fatwā refusal/triage), grounded in classical foundations and a reliable English reference to the Quran.
2. Education-specific operationalization: Two Qur'anic education use cases (Adaptive Tahfiz Tutor; Tafsīr-grounded RAG) with risk–control maps and Arabic/recitation-aware checks, aligned to evidence on education AI and speech/dialect issues.
3. Auditable artefacts out of the box: Lifecycle gates, transparency/model notes, scholar review forms, a Shariah governance checklist, and a versioned provenance ledger, following best practice on model/data documentation and risk management.
4. Measurable assurance, not slogans: Calibration (ECE), refusal/uncertainty protocols, and fairness audits across gender/age/dialect, with incident registers and KPIs
5. Institutional fit and compliance: A maturity model and roles (Shariah board + AI ethics) and a crosswalk to UNESCO/OECD/NIST/EU AI Act for smooth adoption

Table 1 shows, AIMF moves from principle to practice by pairing Qur'anic aims with measurable controls and audit-ready artefacts.

Table 1: Where AIMF differs from existing frameworks (summary of the gap and our additions).

| Common gap | What AIMF adds |
|--------------------------------------|--|
| Principles stop at policy [8, 9, 10] | Values → controls → lifecycle evidence with concrete artefacts (gates, templates, logs) [10, 21, 27, 26]. |
| Generic education focus [6, 1] | Qur’an/tafsīr corpus integrity, qirā’āt awareness, Arabic-first UX, fatwā refusal routes [12, 13, 14, 15]. |
| No source authentication norm | Whitelisted, versioned sources with citations by default and a provenance ledger [26, 27]. |
| Vague “fairness” guidance [9] | Dialect/accnt parity tests for recitation/retrieval; documented remediation plans [38, 24]. |
| Hard to adopt quickly | Ready-to-use checklists, transparency notes, scholar-review forms, and a maturity model tied to audits [10, 21, 27]. |

BACKGROUND AND RELATED WORK

Maqāsid al-Qur’an and the Darūriyyāt

Classical and contemporary scholars articulate higher purposes that guard religion (hifz al-dīn), life (nafs), intellect (‘āql), progeny (nasl), and wealth (māl), with expansions to environment and dignity. Our framework treats these as non-negotiable guardrails and translates them into system requirements for AI in education.

AI in IR4.0 Education

IR4.0 enables adaptive learning, multimodal assessment, ubiquitous access, and data-driven personalization. Yet educational AI raises risks privacy, bias, opacity, hallucinations, over-reliance, and erosion of teacher authority requiring governance that respects Islamic pedagogy and ethics.

Related Governance

International instruments UNESCO’s ethics recommendation, OECD AI Principles, NIST AI RMF and jurisdictional laws (e.g., EU AI Act) converge on risk-based, rights-preserving oversight. AIMF complements these by adding a Qur’an-first layer for intent, scope, and red lines in religious education.

Table 2: Crosswalk: AIMF controls vs. major governance instruments.

| AIMF Control | Mapped Instruments |
|------------------------------------|--|
| Authenticated sources & provenance | NIST AI RMF (MAP/MEASURE); UNESCO (Transparency); OECD (Human-centric) |
| Doctrinal boundaries & refusals | EU AI Act (use-case restrictions); OIC-IPHRC (Ethical guardrails) |
| Calibration & uncertainty | NIST (Measure/Manage); ISO/IEC 23894 (Risk treatment) |
| Child-safe defaults | UNESCO AIED; national child protection laws/policies |
| Privacy by design | GDPR; PDPA (MY); ISO/IEC 23894; IMDA Model AI Framework |
| Fairness audits | OECD AI Principles (Fairness); NIST (Harm mitigation) |
| Green compute | Institutional sustainability policies; reporting norms |

Quranic & Prophetic Foundations

We ground AIMF in four working principles and draw on canonical sources so the implementation stays faithful and practical

- Justice and Non-Discrimination

•

The Qur'an's call to return trusts and judge fairly speaks directly to how we curate data, score learners, and manage access.

God commands you to return trusts to their rightful owners; and when you judge between people, judge with justice.

(Qur'an 4:58).

People, We created you from a male and a female and made you peoples and tribes so that you may know one another. The most noble of you before God is the most mindful of Him.

(Qur'an 49:13).

Help your brother, whether he is wronging or wronged. They asked, "How if he is wronging?" He said, by restraining him from wrongdoing.

(Ṣaḥīḥ al-Bukhārī)

In practice: set and publish a sensitive-attribute policy, audit parity across age/gender/region, and document remediation plans.

- Truthfulness and Epistemic Humility

The charge to speak plainly and avoid conjecture translates into disciplined retrieval and calibrated answers

"Believers, be mindful of God and speak words that are straight and true".

(Qur'an 33:70)

"They have no knowledge of it; they follow only conjecture and conjecture does not avail against the truth"

(Qur'an 53:28)

The signs of a hypocrite are three: when he speaks he lies; when he promises he breaks it; and when entrusted he betrays.

(Ṣaḥīḥ al-Bukhārī; Ṣaḥīḥ Muslim).

In practice: constrain generation to authenticated sources, surface uncertainty bands, and refuse when evidence is thin.

.

- Dignity, Mercy, and Well-Being

Respect for honor, gentle conduct, and the sanctity of life should shape UX, moderation, and safeguarding.

"

Believers, let not one group ridicule another ... do not insult one another nor use offensive nicknames."

(Qur'an 49:11).

"Believers, avoid much suspicion ... and do not backbite one another."

(Qur'an 49:12)

"Whoever kills a person ... it is as if he killed all people; and whoever saves a life, it is as if he saved all people".

(Qur'an 5:32).

God is gentle and loves gentleness in all things.

(Ṣaḥīḥ Muslim)

In practice: child-safe defaults, clear opt-outs, respectful feedback, and firm escalation paths.

- Seeking Knowledge with Responsibility

Seeking knowledge calls for care, sincerity, and accountability in the tools we build. Say, "My Lord, increase me in knowledge."

(Qur'an 20:114).

God raises in rank those of you who believe and those who have been given knowl- edge.

(Qur'an 58:11).

Whoever follows a path seeking knowledge, God makes a path to Paradise easy for him.

(Ṣaḥīḥ Muslim).

In practice: teacher override mechanisms, scholar sign-offs, and clear limits on the system's remit.

Controls Catalogue (Concise)

| Maqāṣid Aim | Control | Operational Evidence |
|-----------------------|---|---|
| Hifz al-dīn | Retrieval only from authenticated muṣḥaf, canonical qirā'āt, vetted tafāsīr | Source whitelist; corpus versioning; citations in UI |
| Hifz al-nafs | Age gates; crisis/abuse escalation; respectful feedback | Incident register; escalation SOP; response SLAs |
| Hifz al-'āql | Calibration; uncertainty notes; refusal on ambiguity | ECE/qual checklist; transparency notes; refusal templates |
| Hifz al-nasl | Parental dashboards; session privacy defaults | Consent artifacts; DPO sign-off; access logs |
| Hifz al-māl | Data minimization; sustainable billing | DPIA; billing transparency report |
| Dignity & environment | Inclusive UX; green compute reporting | Accessibility audit; energy/inference log |



Figure 1: Lifecycle gates for AIMF; each gate requires concrete artefacts and role sign-offs. [10, 21, 26, 27]

THE AI-MAQASID FRAMEWORK (AIMF)

AIMF is a layered scaffold from values → principles → design controls → lifecycle evidence. **Values** →

Principles

- Hifz al-dīn (Preservation of religion): Guard textual integrity; prohibit distortions and unauthorized doctrinal outputs; privilege authentic sources.
- Hifz al-nafs (Life/wellbeing): Avoid harm (psychological, social); ensure accessibility and inclusion for diverse learners, including special needs.
- Hifz al-'āql (Intellect): Promote sound reasoning; surface epistemic status (certainty vs. speculation); avoid hallucinations; calibrate confidence.
- Hifz al-nasl (Progeny/dignity): Child-safe defaults; parental/guardian controls; age-appropriate content; anti-grooming safeguards.
- Hifz al-māl (Wealth/trust): Protect data/fees; fair access; prevent predatory monetization; ensure sustainability.
- Dignity & environment (expanded aims): Respect user dignity, culture, and ecological impact (compute, energy).

Design Controls (selected)

- Source Guardrails: Retrieval only from authenticated Qur'an text, canonical qirā'āt, and vetted tafāsīr; provenance tracking.
- Doctrinal Boundaries: Refusal policies on issuing fatāwā; route fiqh queries to qualified scholars; model cards declaring scope limits.
- Privacy by Design: Minimized data collection, client-side processing where feasible; differential privacy/federated learning for analytics.

- Fairness & Inclusion: Accent-aware recitation assessment; bias testing by gender/age/region; accessibility (screen readers, dyslexia-friendly fonts).
- Transparency: System cards and concise “transparency notes” in learner-facing language; uncertainty communication (calibrated confidence).
- Safety: Abuse/threat detection; age gates; human-in-the-loop escalation; content filters for sanctity of scripture.
- Sustainability: Efficient models/edge inference; green compute reporting.

Lifecycle Evidence (Govern–Map–Measure–Manage)

- Govern: Shariah Governance Board, ethics committee, role definitions, incident response and red-team protocols.
- Map: Risk register, data lineage, permissible use catalog, stakeholder mapping (teachers, parents, huffāz, asātizah, learners, regulators).
- Measure: Bias & calibration audits, privacy tests, reliability metrics, drift monitoring, periodic Shariah compliance reviews.
- Manage: Mitigation plans, model updates, user communication, rollback procedures, continuous improvement.
-

Reference Architecture

Figure 2 shows how governance (Shariah Board and AI governance) constrains sources and policy, which in turn govern retrieval/provenance and downstream AI services. The assurance layer (bias, calibration, red-teaming) feeds results back into services and audit logs, while learner/teacher UIs present outputs. Arrows indicate control flow and oversight loops across the stack. Read top–down: governance sets remit and red lines; authenticated sources and policy gate what enters the system; retrieval anchors answers in vetted texts; services implement pedagogy; assurance closes the loop with measurement and mitigation; the UI and audit trail externalize transparency and accountability

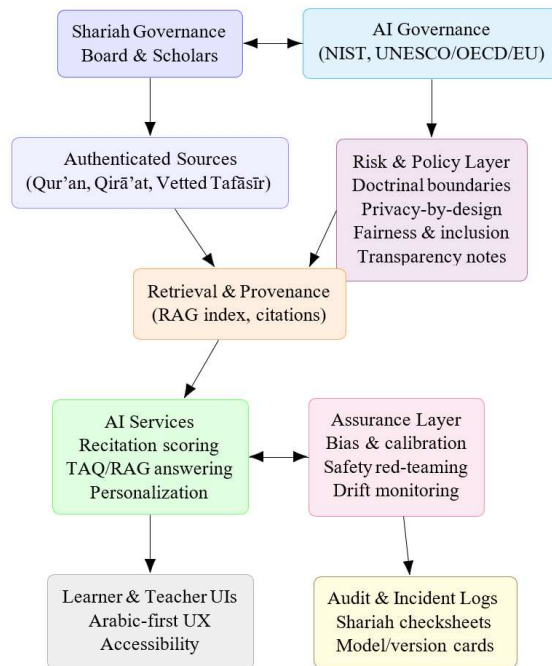


Figure 2: AI-Maqasid reference architecture (conceptual).

Figure 2: AI-Maqasid reference architecture (conceptual).

METHODS: FROM VALUES TO CONTROLS

We adopt a design-science and participatory approach: (1) derive requirements from Maqāsid and teaching practice; (2) co-design controls with educators, ‘ulama, technologists; (3) implement minimally invasive prototypes; (4) evaluate with mixed methods (metrics + scholar review); (5) institutionalize via policy and training.

USE CASES

1. Adaptive Tahfiz Tutor

Goal: Improve memorization (hifz) and tajwīd with personalization that preserves sanctity.

Data: On-device audio (default), optional secure cloud; consent-driven analytics.

Features: Accent-aware recitation scoring, tajwīd feedback, spaced repetition plans, respectful nudges.

AIMF Controls:

- Hifz al-dīn: Correct only to canonical recitations; display verses with authoritative muṣḥaf layout; cite riwāyah.
- Privacy: Client-side inference where possible; aggregated metrics with differential privacy; parental dashboards for minors.
- Fairness: Evaluate scoring parity by accent/age/gender; calibrate confidence; allow teacher overrides.
- Transparency: Show how scores are computed (high-level), with confidence bands and references.

2. Quranic Studies RAG Assistant (Tafsir-Grounded)

Goal: Answer questions with citations to authenticated sources; refuse fatwa/out-of-scope. Pipeline: Retrieval limited to vetted corpora (Qur’an text, recognized tafsāsīr, peer-reviewed literature); provenance links for every claim.

AIMF Controls:

- Doctrinal boundaries: Built-in refusal policies; route fiqh queries to scholars; warnings on speculative content.
- Safety: Prevent disrespectful prompts; log/report abuse; moderation tuned for sanctity of scripture.
- Explainability: Always show sources; highlight exact verses/tafsīr passages; expose uncertainty.

Risk–Control Map for the Two Use Cases

Table 4: Risks and bound controls for the Adaptive Tahfiz Tutor and Tafsīr–Grounded RAG Assistant.

| Risk | Tahfiz Tutor Controls | RAG Assistant Controls |
|--------------------------------------|--|---|
| Hallucination/ Incorrect Guidance | On–device constrained prompts; show canonical text; teacher override | Retrieval from vetted tafsīr; always cite; refuse out–of–scope fiqh |
| Bias/Exclusion | Accent–aware scoring; parity audits | Query–language normalization; review parity across dialects |
| Privacy | Client–side inference; DP aggregates | No retention by default; hashed logs; DPIA |
| Child Safety | Age gates; guardian dashboards | Harms filter; escalation SOP |
| Over–reliance | Study plans with reflection prompts | Uncertainty banners; “ask a scholar” routes |

EVALUATION: METRICS AND CHECKLISTS

Table 5: Mapping Maqāsid to design controls, metrics, and evidence.

| Maqāsid Aim | Primary Controls | Quant/Qual Metrics | Evidence Artifacts |
|-----------------------|--|--|--|
| <i>Hifẓ al-dīn</i> | Authenticated sources; doctrinal boundaries; refusal flows | % answers with citations; doctrinal review pass-rate; misuse incidents | Scholar review log; source whitelist; refusal templates [27] |
| <i>Hifẓ al-nafs</i> | Safety filters; age gates; well-being nudges | Abuse detection precision/recall; incident rate; user well-being surveys | Incident register; red-team reports [10] |
| <i>Hifẓ al-‘āql</i> | Calibration; uncertainty communication | Expected Calibration Error (ECE); accuracy@top- <i>k</i> ; teacher override rate | Model card; calibration plots [32, 27] |
| <i>Hifẓ al-nasl</i> | Parental controls; child-safe defaults | Minor user protection tests; COPPA/PDPA checks | DPO reports; consent flows [39, 40] |
| <i>Hifẓ al-māl</i> | Data minimization; secure billing | Breach count; time-to-remediation; fee transparency score | DPIA; audit reports [41] |
| Dignity & environment | Inclusive UX; green compute | Accessibility conformance; energy per inference | Accessibility audit; sustainability note [42, 28] |

Expert Review Rubric (No Experiments Required)

Table 6: Structured rubric for scholar/educator review (each 1–5; median reported).

| Criterion | Who Rates | Target (5) |
|------------------------------|-------------------|------------|
| Textual fidelity (sources) | Scholar panel | 5.0 |
| Doctrinal boundary handling | Scholar panel | ≥4.5 |
| Clarity & transparency notes | Educators | ≥4.5 |
| Child–safety posture | Safeguarding lead | ≥4.5 |
| Fairness & inclusion | QA/Audit | ≥4.0 |
| Privacy posture | DPO | ≥4.5 |
| Operational readiness | Program lead | ≥4.0 |

Reporting. Publish a one–page summary with medians, free–text concerns, and remediation actions. Keep signed checklists as audit artefacts.

IMPLEMENTATION ROADMAP (GLOBAL)

institutional governance: Create/empower a Shariah Governance Board and an AI Ethics Committee with clear charters and escalation paths.

1. Policy alignment: Crosswalk AIMF with national policy (e.g., PDPA, AI governance guidelines) and international frameworks (UNESCO, OECD, NIST, EU AI Act)
2. Curricular integration: Teacher training on AI literacy, Maqāṣid-based evaluation, and transparency note-writing.
3. Technical stack: Privacy-first data architecture; retrieval-only from authenticated sources; continuous bias and calibration audits
4. Community engagement: Parents, asātizah, tahfiz institutions, NGOs; publish annual assurance reports.
5. Institutional Maturity Model for AIMF

Table 7: AIMF maturity levels and expected artefacts.

| Level | Characteristics | Required Artefacts |
|-----------------------|---|--|
| L1: Initial | Ad hoc pilots; undocumented boundaries | Scope note; minimal corpus whitelist |
| L2: Managed | Named roles; periodic reviews | DPIA; transparency notes; refusal templates |
| L3: Defined | RACI; lifecycle gates enforced | Risk register; parity report; calibration note |
| L4: Quantified | KPIs tracked; green compute | ECE trend; incident MTTR; energy/inference [32, 28] |
| L5: Optimizing | Continuous improvement; public reporting | Annual assurance report; versioned corpus ledger |

DISCUSSION

Benefits. AIMF anchors AI-enabled pedagogy in Qur'anic purposes while leveraging modern assurance (risk registers, audits). It respects teacher authority and student dignity, reduces hallucinations via constrained retrieval, and makes governance inspectable. Risks & Limits. Over-reliance on automation, residual bias in speech models across dialects, and context loss in short answers remain challenges. Scholarly disagreement across madhāhib requires clearly declared scope and versioned corpora. Continuous oversight is essential.

CONCLUSION

AIMF offers a practical path to bring Qur'anic ethics into AI design for education in the IR4.0 era. By translating values into controls and evidence, institutions can adopt AI responsibly, protect learners and scripture, and accelerate high-quality, inclusive Quranic teaching.

REFERENCES

- O. Zawacki-Richter, V. I. Marín, M. Bond, and F. Gouverneur, "Systematic review of the application of artificial intelligence in higher education: where, how and to what extent?," *International Journal of Educational Technology in Higher Education*, 16(39), 2019. doi: 10.1186/s41239-019-0172-0.
- A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," *Nature Machine Intelligence*, 1, 389–399, 2019. doi: 10.1038/s42256-019-0088-2.
- Y. I. S. al-Nawawi, *Al-Arba'in al-Nawawiyya (Forty Hadith)*. Critical Arabic editions; reliable English translations available.
- M. J. al-Tabari, *Jami' al-Bayan 'an Ta'wil Ay al-Qur'an*. Beirut: Dar al-Fikr, n.d.
- M. A. al-Qurtubi, *Al-Jami' li-Ahkam al-Qur'an*. Beirut: Dar al-Kutub al-'Ilmiyyah, n.d.
- UNESCO, *AI and Education: Guidance for Policy-makers*. Paris: UNESCO, 2021. ISBN 978-92-3-100447-1. Available: <https://unesdoc.unesco.org/ark:/48223/pf0000376709>
- K. Schwab, *The Fourth Industrial Revolution*. Geneva: World Economic Forum, 2016.
- UNESCO, *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO, 2021. Available: <https://unesdoc.unesco.org/ark:/48223/pf0000381137>
- OECD, *Recommendation of the Council on Artificial Intelligence*. Paris: OECD, 2019. OECD/LEGAL/0449. Available: <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- National Institute of Standards and Technology (NIST), *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1, 2023. doi: 10.6028/NIST.AI.100-1.
- European Union, *Artificial Intelligence Act (Regulation (EU) 2024/1689 of 13 June 2024)*. Official Journal of the European Union, L202, 12 July 2024.
- The Qur'an, trans. M. A. S. Abdel Haleem. Oxford: Oxford University Press, 2004.
- A. I. al-Shāṭibī, *Al-Muwāfaqāt fī Uṣūl al-Sharī'ah*. Beirut/Cairo scholarly editions.
- M. al-Ṭāhir Ibn 'Āshūr, *Treatise on Maqasid al-Shariah*, trans. M. E. T. El-Mesawi. Herndon, VA: IIIT, 2006.
- J. Auda, *Maqasid al-Shariah as Philosophy of Islamic Law: A Systems Approach*. Herndon, VA: IIIT, 2008.

- W. Holmes, N. Bialik, and C. F. Adel, “Ethics of AI in Education: Towards a Community- Wide Framework,” *International Journal of Artificial Intelligence in Education*, 32, 2022. doi: 10.1007/s40593-021-00239-1.
- Z. Ji, N. Lee, R. Frieske, T. Yu, et al., “Survey of Hallucination in Natural Language Gen- eration,” *ACM Computing Surveys*, 55(12), 2023. doi: 10.1145/3571730.
- C. Mutimukwe, O. Viberg, L.-M. Öberg, and T. Cerratto-Pargman, “Students’ privacy concerns in learning analytics: Model development,” *British Journal of Educational Tech- nology*, 53(6), 2022. doi: 10.1111/bjet.13234.
- K. M. L. Jones, “Learning analytics and higher education: a proposed model for establish- ing informed consent,” *International Journal of Educational Technology in Higher Edu- cation*, 16(24), 2019. doi: 10.1186/s41239-019-0155-0.
- ISO/IEC 23894:2023, Artificial Intelligence — Risk Management. Geneva: ISO/IEC, 2023.
- OIC–IPHRC, Jeddah Declaration on Ethical Challenges of Artificial Intelligence, 2024. Available: <https://www.oic-iphcr.org/>
- Infocomm Media Development Authority (IMDA), Singapore, Model AI Governance Framework (2nd ed.), 2020.
- N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, “A survey on bias and fairness in machine learning,” *ACM Computing Surveys*, 54(6):115, 2021. doi: 10.1145/3457607.
- L. Floridi, J. Cowls, M. Beltrametti, R. Chatila, et al., “AI4People—An ethical framework for a good AI society,” *Minds and Machines*, 28, 689–707, 2018. doi: 10.1007/s11023- 018-9482-5.
- T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. M. Wallach, H. Daumé III, and K. Crawford, “Datasheets for Datasets,” *Communications of the ACM*, 64(12), 86–92, 2021. doi: 10.1145/3458723.
- M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru, “Model Cards for Model Reporting,” in *FAT* 2019*, pp. 220–229. doi: 10.1145/3287560.3287596.
- R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, “Green AI,” *Communications of the ACM*, 63(12), 54–63, 2020. doi: 10.1145/3381831.
- C. Dwork, “Differential Privacy,” in *ICALP 2006*, LNCS 4052, Springer, 2006, pp. 1–12. doi: 10.1007/117870061.
- H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Aguëra y Arcas, “Communication-Efficient Learning of Deep Networks from Decentralized Data,” *AIS- TATS 2017*, PMLR 54:1273–1282, 2017.
- J. Buolamwini and T. Gebru, “Gender Shades: Intersectional Accuracy Disparities in Com- mercial Gender Classification,” *FAT* 2018*, PMLR 81:77–91, 2018.
- C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On Calibration of Modern Neural Net- works,” *ICML 2017*, PMLR 70:1321–1330, 2017.

- A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research," *MIS Quarterly*, 28(1):75–105, 2004.
- K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A Design Science Research Methodology for Information Systems Research," *Journal of Management Information Systems*, 24(3):45–77, 2007.
- E. Björgvinsson, P. Ehn, and P.-A. Hillgren, "Design Things and Design Thinking," *CoDesign*, 8(2–3):101–116, 2012.
- R. B. Johnson and A. J. Onwuegbuzie, "Mixed Methods Research: A Research Paradigm Whose Time Has Come," *Educational Researcher*, 33(7):14–26, 2004.
- N. J. Cepeda, E. Vul, D. Rohrer, J. T. Wixted, and H. Pashler, "Spacing Effects in Learning: A Temporal Ridge of Optimal Retention," *Psychological Science*, 19(11):1095–1102, 2008.
- A. Koenecke, A. Nam, E. Lake, J. Nudell, M. Quartey, Z. Mengesha, C. Toups, J. R. Rickford, D. Jurafsky, and S. Goel, "Racial disparities in automated speech recognition," *PNAS*, 117(14):7684–7689, 2020. doi: 10.1073/pnas.1915768117.
- United States Federal Trade Commission, "Children's Online Privacy Protection Rule (COPPA)," 16 CFR Part 312.
- [Government of Malaysia, Personal Data Protection Act 2010 (PDPA) and subsidiary regulations, 2010.
- European Union, General Data Protection Regulation (GDPR), Regulation (EU) 2016/679, 2016.
- W3C, Web Content Accessibility Guidelines (WCAG) 2.2, 2023. Available: <https://www.w3.org/TR/WCAG22/>
- M. I. al-Bukhārī, Ṣaḥīḥ al-Bukhārī. Riyadh: Darussalam, 1997.
- [Muslim ibn al-Ḥajjāj, Ṣaḥīḥ Muslim. Riyadh: Darussalam, 2007.
- [Abu Hāmid al-Ghazālī, *Al-Mustasfa min 'Ilm al-Uṣul*. Beirut: Dar al-Kutub al-'Ilmiyyah, n.d.
- F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," arXiv:1702.08608, 2017.