Human Factors, Architecture, Sustainable Urban Planning and Infrastructure, Vol. 229, 2026, 220–230
https://doi.org/10.54941/ahfe1007970

Human-Level Barriers to Generative AI Adoption in Architectural Practice

Silvia Albano

Xi'an Jiaotong - Liverpool University, Suzhou, Jiangsu 215123, China

ABSTRACT

Generative artificial intelligence is entering architectural practice rapidly, although it is not yet possible to speak of full adoption, since it is used for informal experimentation such as ideation, visualisation, drafting, text production, and workflow acceleration. This narrow interpretation overlooks a deeper human-level adoption gap. The central challenge is not only whether AI systems can produce useful design-related outputs, but whether architectural practices can integrate them responsibly within workflows shaped by key professional and procedural constraints. This paper addresses this gap by examining the barriers, readiness conditions, and governance requirements for responsible generative AI adoption in architectural practice. The study is a PRISMA-guided systematic review and thematic synthesis of peer-reviewed articles, proceedings, book chapters, and selected governance documents (2014–2026) from architecture, design research, HCI, human factors, automation, AI ethics, and governance. Instead of evaluating individual AI tools or collaboration models, it analyses reported human-level barriers and enablers across five interrelated dimensions: trust and reliance, accountability and provenance, skills and role transformation, organisational readiness, and professional governance. The synthesis is then mapped to RIBA project stages to identify where risks, responsibilities, and verification thresholds change across the project lifecycle. The paper makes three contributions: a taxonomy of human-level barriers to generative AI adoption, a stage-specific risk and responsibility map, and a practical adoption-readiness checklist for architectural firms. The paper argues that responsible AI adoption in architecture is a human and organisational challenge: AI can support professional work only when embedded within transparent, stage-gated, and accountable systems of judgement.

Keywords: Generative AI, human-AI interaction, architectural practice professional judgement, responsible AI

INTRODUCTION

Generative artificial intelligence is moving rapidly from experimental novelty to everyday design support in architectural practice. It is now used for precedent search, image generation, concept communication, meeting summaries, specification drafting and early regulatory prompts. Professional evidence confirms this acceleration: the RIBA Artificial Intelligence Report 2025 reports that 59% of responding architectural practices use AI, compared with 41% in 2024, while also stressing that AI is expected to augment rather than replace the architect's role (Royal Institute of British Architects, 2025). This shift makes adoption a live professional issue rather than a future speculation.

Received May 17, 2026; Revised May 26, 2026; Accepted June 11, 2026; Available online July 20, 2026

At the same time, the research literature shows that architectural AI use remains concentrated in early-stage design, image-based workflows and visual/material exploration. Jang, Roh and Lee's PRISMA-based review of 161 journal papers found that AI use is concentrated in schematic design phases, with images dominating input and output data and subjective assessment still central to evaluation (Jang et al., 2025). Paananen, Oppenlaender and Visuri's laboratory study with architecture students also shows that text-to-image systems can support serendipitous ideation, but only when users understand design constraints and treat outputs as prompts for critique rather than resolved proposals (Paananen et al., 2024). These findings suggest an adoption paradox: AI is increasingly accessible, but professional reliability, evidence and responsibility do not automatically follow.

This paper argues that the main barrier to generative AI adoption in architecture is not tool availability but human-level and organisational readiness. Architectural decisions affect spatial quality, safety, accessibility, cost, environmental performance, legal compliance and public value. Therefore, AI outputs can enter professional workflows only through systems of critical judgement, verification, confidentiality control and accountable sign-off. The paper asks: RQ1: Which human-level barriers and enablers shape the adoption of generative AI in architectural practice? RQ2: How do these barriers and enablers vary by task type and project stage? RQ3: Which governance measures can support responsible, scalable and professionally accountable AI adoption within architectural firms?

The contribution is threefold. First, the paper develops a taxonomy of human-level adoption barriers. Second, it maps risks and responsibilities across architectural project stages, using the RIBA Plan of Work as a practical reference structure (Royal Institute of British Architects, 2020). Third, it proposes a stage-gated adoption-readiness framework and checklist for architectural firms. The argument deliberately avoids framing the paper as another human-AI collaboration model. Instead, it examines the readiness conditions that allow generative AI to be used responsibly in professional practice.

LITERATURE REVIEW

Generative AI and Architectural Adoption

AI in architecture builds on earlier computational design traditions, including parametric modelling, optimisation, simulation and rule-based generation. Generative AI changes the adoption landscape because it is accessible through natural language prompts and general-purpose platforms. This lowers the threshold for non-specialist use and allows AI to enter everyday practice through image generation, text generation, summarisation, coding assistance and documentation support.

Recent reviews confirm both the breadth and immaturity of this field. Li et al. (2025) and Horvath and Pouliou (2024) map generative AI applications across design steps and conceptual exploration, while Khan, Chang and Chang (2025) argue that current tools risk repeating early CAD limitations if they remain isolated task-automation systems rather than integrated, multidisciplinary workflows. Onatayo, Abanda and Oyedele (2024) similarly show that generative AI applications in AECO extend from design and BIM to construction and operations, but adoption is still fragmented. These studies justify shifting attention from technical capability to adoption readiness.

The adoption gap appears when AI outputs move from private experimentation into project workflows. A generated image may influence client expectations before feasibility is tested; a generated report paragraph may appear authoritative while lacking verified sources; and a generated compliance summary may omit local regulatory nuance. The question is therefore not simply what AI can generate, but under what conditions AI-generated outputs can become part of accountable architectural work.

Trust, Automation Bias and Human-AI Interaction

Human factors literature provides a necessary basis for understanding this gap. Lee and See (2004) frame trust in automation as a decision to rely on an automated system under uncertainty. Parasuraman and Riley (1997) distinguish between use, misuse, disuse and abuse of automation, while Parasuraman and Manzey (2010) show how complacency and automation bias arise when users monitor automated outputs less critically. In architectural practice, this becomes critical because AI outputs can look visually complete and linguistically authoritative even when they are technically, legally or spatially weak.

Recent human-AI research expands this concern. Bansal et al. (2019) show that performance depends on users' mental models of when AI succeeds or fails, while Bucinca, Malaya and Gajos (2021) show that cognitive forcing can reduce over-reliance by requiring users to pause before accepting AI advice. Glikson and Woolley (2020) further show that trust depends on task, system behaviour and interpretation. For architectural practice, this supports review moments wherever overconfidence is likely.

Human-AI interaction guidelines reinforce this point. Amershi et al. (2019) emphasise communicating capability, supporting correction, showing uncertainty and enabling user control. Explainability must also be meaningful to the user's task (Miller, 2019): for architects, explanation must help judge feasibility, context, compliance, accessibility and evidence.

Accountability, Authorship and Professional Governance

Generative AI complicates accountability because outputs may combine prompts, training data, system behaviour and human editing. Architectural outputs are already collective, but AI adds opacity to the chain of contribution. This has implications for authorship, intellectual property, client communication, quality assurance, professional liability and academic integrity. The AIA Trust (2025) links generative AI directly to professional obligations such as competence, candour, confidentiality, proper attribution and supervisory responsibility, noting that architects must remain in responsible control of work they sign or seal.

AI accountability literature reinforces the need for traceability. Raji et al. (2020) argue for internal auditing to close accountability gaps. Although architectural firms may not train foundation models, they still need equivalent use-side documentation: which tool was used, what data were entered, what output was generated, who edited it, which sources were checked and who approved the final decision.

Governance frameworks provide broader reference points. NIST AI RMF organises AI risk around governance, mapping, measurement and management (National Institute of Standards and Technology, 2023). ISO/IEC 42001 treats AI as an organisational management-system issue (ISO/IEC, 2023), while the EU AI Act stresses risk classification, transparency, documentation and human oversight (European Parliament and Council of the European Union, 2024). Legal and AEC ethics research further links AI adoption to privacy, IP, cybersecurity, liability, transparency and responsibility (Liang et al., 2024; Novelli et al., 2024).

Skills, Professional Learning and Role Transformation

AI adoption also affects learning and role definition. Architectural judgement develops through drawing, modelling, critique, revision, coordination, site exposure and technical review (Schön, 1983; Cross, 2006). Many tasks that AI accelerates - precedent search, visualisation, report drafting, diagramming and option comparison - are also learning tasks through which junior staff acquire tacit knowledge.

This tension confirms the paper's central claim: adoption readiness is not only technical capacity but professional culture. Firms must train staff to recognise hallucination, identify data sensitivity, disclose AI contribution, verify sources and distinguish speculative generation from professional validation.

METHODOLOGY AND PRISMA-GUIDED EVIDENCE BASE

The study is a PRISMA-guided systematic literature review and thematic synthesis. PRISMA 2020 supports transparency in search, screening and inclusion (Page et al., 2021). The synthesis is qualitative rather than statistical: the unit of analysis is the reported or implied adoption barrier, adoption enabler, professional risk or governance mechanism.

The review covers literature from 2014 to 2026 across architecture, design research, HCI, human factors, automation, AI ethics, AECO decision-support and governance. Sources were included when they addressed AI-supported design or AECO decision-making; human-level factors such as trust, automation bias, skills, authorship or accountability; or governance issues such as confidentiality, liability, auditability and data protection. Technical papers focused only on model performance were excluded unless they addressed workflow, evaluation, human use or professional implications.

To remain within a short conference format, the PRISMA material is reported compactly. The review does not replace larger systematic reviews such as Jang et al. (2025); it uses them as evidence anchors and extends them through human-factors and governance literature. Data extraction proceeded in two passes: first by domain, AI application, task and stage; second by barrier, enabler, risk level and governance implication.

Table 1: PRISMA-guided evidence base and coding role.

Evidence Group	Reference-Supported Data Integrated in the Review	Contribution to Coding
Architecture and generative AI	Jang et al. (2025) reviewed 161 papers; Li et al. (2025), Horvath and Pouliou (2024) and Paananen et al. (2024) frame design-stage use and ideation constraints.	Application area, design stage, evaluation weakness and early-stage overconcentration.
AEC and practice integration	Onatayo et al. (2024) and Khan et al. (2025) identify fragmented workflow integration and limited evaluation standards.	Workflow integration, scalability and practice-readiness barriers.
Human factors and HAI	Lee and See (2004), Parasuraman and Riley (1997), Parasuraman and Manzey (2010), Bansal et al. (2019), Bucinca et al. (2021), Amershi et al. (2019), Miller (2019) and Glikson and Woolley (2020) explain trust, over-reliance, mental models and uncertainty.	Trust calibration, automation bias, correction, explanation and review thresholds.
Governance and professional responsibility	AIA Trust (2025), NIST (2023), ISO/IEC (2023), EU AI Act (2024), Raji et al. (2020), Liang et al. (2024), Novelli et al. (2024) and RIBA (2025) address competence, confidentiality, auditability and risk governance.	Accountability, provenance, data control, disclosure, sign-off and organisational readiness.

FINDINGS

Human-Level Barriers and Readiness Responses

The synthesis identifies eight interrelated barriers. Trust calibration concerns the ability to judge when AI outputs are reliable, uncertain or unsuitable. Automation bias and visual overconfidence arise when fluent text or persuasive images are interpreted as decision-ready before feasibility is tested. Accountability ambiguity concerns responsibility for prompts, outputs, edits, verification and sign-off. Authorship and provenance uncertainty affect attribution, quality assurance and professional liability. Skills gaps concern AI literacy, source checking and recognition of hallucination. Role ambiguity concerns changing responsibilities for junior staff, project architects, BIM coordinators, consultants and directors. Confidentiality risks arise when client or project information enters unapproved AI tools. Governance immaturity refers to the absence of policies, review thresholds, escalation rules and project learning routines.

The barriers reinforce one another. A lack of AI literacy increases automation bias; weak provenance increases accountability ambiguity; and unclear roles make governance difficult to operationalise. This means adoption readiness cannot be achieved by a single policy or training session. It requires a combined system of professional judgement, documented decision-making, role clarity and organisational learning.

Table 2: Human-level barriers and corresponding readiness responses.

Barrier	Architectural Risk	Readiness Response
Trust calibration	Outputs are over-trusted or under-used because capability is misunderstood.	Classify tasks by risk and verify outputs according to stage.
Automation bias / visual overconfidence	Images or text appear resolved before feasibility, compliance or performance testing.	Use critique, cognitive forcing and explicit uncertainty notes.
Accountability ambiguity	Responsibility for AI-supported decisions becomes unclear across teams.	Define reviewer, approver and sign-off responsibility.
Authorship / provenance uncertainty	AI contribution is invisible in reports, images, diagrams or specifications.	Record tool, input, output, edits, checks and final decision.
Skills and AI literacy gaps	Users may not detect hallucination, bias or weak evidence.	Teach verification, data sensitivity and professional limits.
Role ambiguity	AI changes task allocation and may weaken learning pathways.	Define who may use AI, who reviews it and how learning is preserved.
Confidentiality and IP risk	Project or client data may enter unsafe platforms.	Use approved tools and client-sensitive information controls.
Governance immaturity	AI remains individual, informal and unscalable.	Develop firm policy, review thresholds, logs and post-project learning.

Stage-Specific Risk and Responsibility Map

AI-related risk changes across project stages. Early-stage uses are usually exploratory, but they can still distort expectations if clients or students interpret images as feasible proposals. Later-stage uses carry stronger technical, contractual and safety implications. The RIBA Plan of Work is therefore used as a practical stage reference, not as a new design model. It helps clarify where review, escalation and evidence requirements intensify.

Table 3: Stage-specific AI risk and professional control.

Project Stage	Typical AI Use	Human-Level Risk	Required Control
Strategic Definition	Market scans; precedent summaries; strategic scenarios.	Generic assumptions; weak evidence; premature framing.	Source checking, uncertainty notes and senior review.
Preparation and Briefing	Stakeholder maps; user-needs summaries; brief drafts.	Misread client priorities; confidentiality leakage.	Client validation, data controls and brief review.
Concept Design	Visual ideation; narratives; option generation.	Visual overconfidence; authorship ambiguity.	Design critique, feasibility check and AI-use record.
Spatial Coordination	Option comparison; coordination prompts; model-check summaries.	Incomplete assumptions; coordination gaps.	Consultant review, model checking and issue tracking.
Technical Design	Specification drafts; compliance summaries; technical descriptions.	Hallucinated regulations; liability exposure.	Source verification, senior sign-off and no unsupported claims.
Manufacturing and Construction	RFI drafts; sequencing summaries; site issue interpretation.	Contractual and safety risk; inaccurate summarisation.	Controlled use only, formal approval and retained human responsibility.
Handover and Use	O&M summaries; lessons learned; POE and operational feedback.	Missing safety information; privacy risk; biased interpretation.	Document audit, anonymisation and stakeholder review.

Stage-Gated Human-Level AI Adoption Framework

The findings support a Stage-Gated Human-Level AI Adoption Framework with four layers: human judgement, task-risk classification, governance and organisational learning. These layers define when professional judgement must dominate; whether an AI use is low, medium, high or critical risk; which review, escalation, logging and sign-off mechanisms apply; and how project learning is converted into improved protocols.

The framework is deliberately conservative. It does not make AI autonomous in professional decision-making; it identifies where AI can support search, synthesis, generation, comparison or drafting while responsibility remains with competent professionals. Before using AI output, firms should ask: what is the task, what is the risk, what evidence is required, and who reviews and signs off?

Table 4: Task-risk classification and adoption-readiness controls.

Risk Level	Examples	Minimum Controls and Readiness Actions
Low	Brainstorming; non-confidential visual references; internal keyword generation.	User judgement, informal peer check and no confidential data.
Medium	Draft narratives; internal summaries; option-comparison tables; generic precedent scans.	Peer review, source checking where relevant and visible AI-use note.
High	Client-facing text; planning summaries; specification drafts; compliance prompts; cost assumptions.	Senior review, source verification, decision logging and provenance record.
Critical	Health and safety advice; contractual instruction; regulatory submission; liability-sensitive advice.	Specialist review, formal sign-off, escalation to qualified professionals and no autonomous AI decision-making.
Firm-level implementation	Permitted/restricted/prohibited uses; approved tools; data rules; training; post-project review.	Clarify roles, protect client data, preserve junior learning and update protocols through project learning.

Implementation Maturity and Evidence Integration

The framework treats AI readiness as a maturity process rather than a binary condition. A basic firm may allow non-confidential exploratory use; an intermediate firm may add approved tools, guidance and peer review; an advanced firm may embed AI in quality management through logs, provenance records, source verification, staff training and periodic audit. This aligns with evidence that trust depends on mental models, explanation quality and recognition of system failure (Bansal et al., 2019; Glikson & Woolley, 2020; Miller, 2019).

Implementation maturity also depends on the legal and ethical status of outputs. A low-risk internal image may require critique, while AI-supported technical text, regulatory interpretation or contractual communication requires formal verification. Firms should therefore classify outputs by evidential value, data sensitivity and professional consequence, not treat all AI use as equally harmless or unacceptable.

The framework remains adoption-oriented rather than technologically deterministic. Even if systems improve, firms must still decide who may use them, what data can be shared, how outputs are checked, how AI contribution is disclosed and how professional responsibility is retained.

DISCUSSION

The findings show that responsible AI adoption requires a shift from individual experimentation to organisational readiness. This does not reject innovation; it makes innovation trustworthy. RIBA survey data indicate that AI use is becoming normal in practice, while also highlighting uncertainty about risk, role transformation and future practice (Royal Institute of British Architects, 2025). The challenge is therefore cultural and procedural as much as technical.

Professional judgement is stage-specific. During briefing, it concerns client needs and strategic assumptions; during concept design, spatial and experiential quality; during technical design, regulations, accessibility, performance and buildability; and during construction and handover, contractual clarity, site communication and safety. RIBA's Health and Safety Guide reinforces that design responsibility includes risk awareness and safe decision-making (Royal Institute of British Architects, 2023).

The learning implications are also significant. AI may accelerate tasks through which junior staff develop tacit knowledge. It should therefore be used through supervised verification: explaining prompts, identifying assumptions, checking sources, comparing outputs with site and regulatory constraints, and documenting decisions. This links AI literacy to reflective practice and designerly ways of knowing rather than replacing them (Schön, 1983; Cross, 2006).

Governance should be framed as professional infrastructure. As firms already rely on drawing standards, BIM protocols, quality systems and document control, AI governance should make consequential AI-supported decisions visible, reviewable and accountable without documenting every minor prompt. This is vital when firms use general-purpose systems they did not develop and cannot fully inspect.

CONCLUSION AND FUTURE WORK

This paper examined the human-level barriers to generative AI adoption in architectural practice and proposed a governance-oriented approach to responsible implementation. The central argument is that the main challenge is not the availability of AI tools, but the readiness of architectural organisations to use them in ways that are accountable, transparent, stage-appropriate and professionally governed.

The paper contributes three outputs: a taxonomy of human-level adoption barriers, a stage-specific risk and responsibility map, and a Stage-Gated Human-Level AI Adoption Framework supported by an adoption-readiness checklist. The framework is intended as a practical synthesis rather than an empirically validated model. Its value lies in translating human-factors, architectural and governance literature into a compact set of practice-oriented controls.

The study is limited by its literature-based design. Future work should test the framework through interviews, firm case studies and workshops with architects, project managers, BIM coordinators, legal advisors, clients and junior staff. Further research should examine impacts on learning, documentation, insurance, liability, client communication and organisational knowledge. Practical outputs could include AI-use logs, provenance statements, task-risk classification and stage-specific governance templates. Responsible AI adoption will require not only new tools, but critical judgement, shared protocols and organisational learning.

REFERENCES

AIA Trust. (2025). Ethical challenges of generative AI in architectural practice. The American Institute of Architects Trust.

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 1–13. https://doi.org/10.1145/3290605.3300233

Bansal, G., Nushi, B., Kamar, E., Lasecki, W. S., Weld, D. S., & Horvitz, E. (2019). Beyond accuracy: The role of mental models in human-AI team performance. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 1–11. https://doi.org/10.1145/3290605.3300792

Bucinca, Z., Malaya, M. B., & Gajos, K. Z. (2021). To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. Proceedings of the ACM on Human-Computer Interaction, 5(CSCW1), Article 188. https://doi.org/10.1145/3449287

Cross, N. (2006). Designerly ways of knowing. Springer.

European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Official Journal of the European Union.

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. Academy of Management Annals, 14(2), 627–660. https://doi.org/10.5465/annals.2018.0057

Horvath, A. S., & Pouliou, P. (2024). AI for conceptual architecture: Reflections on designing with text-to-text, text-to-image, and image-to-image generators. Frontiers of Architectural Research, 13(3), 593–612. https://doi.org/10.1016/j.foar.2024.02.006

ISO/IEC. (2023). ISO/IEC 42001:2023 Artificial intelligence - Management system. International Organization for Standardization.

Jang, S., Roh, H., & Lee, G. (2025). Generative AI in architectural design: Application, data, and evaluation methods. Automation in Construction, 174, 106174. https://doi.org/10.1016/j.autcon.2025.106174

Khan, A., Chang, S., & Chang, H. (2025). Generative AI approaches for architectural design automation. Automation in Construction, 180, 106506. https://doi.org/10.1016/j.autcon.2025.106506

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Li, C., Zhang, T., Du, X., Zhang, Y., & Xie, H. (2025). Generative AI models for different steps in architectural design: A literature review. Frontiers of Architectural Research. https://doi.org/10.1016/j.foar.2024.12.006

Liang, C.-J., Wang, X., Kassem, M., & Li, H. (2024). Ethics of artificial intelligence and robotics in the architecture, engineering, and construction industry. Automation in Construction, 162, 105369. https://doi.org/10.1016/j.autcon.2024.105369

Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. Artificial Intelligence, 267, 1–38. https://doi.org/10.1016/j.artint.2018.07.007

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

Novelli, C., Casolari, F., Rotolo, A., Taddeo, M., & Floridi, L. (2024). Generative AI in EU law: Liability, privacy, intellectual property, and cybersecurity. Computer Law & Security Review, 55, 106066. https://doi.org/10.1016/j.clsr.2024.106066

Onatayo, D., Abanda, F. H., & Oyedele, L. (2024). Generative AI applications in architecture, engineering, and construction: A systematic review. AI, 4(4), 1076-1100. https://doi.org/10.3390/ai4040046

Paananen, V., Oppenlaender, J., & Visuri, A. (2024). Using text-to-image generation for architectural design ideation. International Journal of Architectural Computing, 22(3). https://doi.org/10.1177/14780771231222783

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hrobjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., McGuinness, L. A., Stewart, L. A., Thomas, J., Tricco, A. C., Welch, V. A., Whiting, P., & Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. Human Factors, 52(3), 381–410. https://doi.org/10.1177/0018720810376055

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), 230-253. https://doi.org/10.1518/001872097778543886

Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 33–44. https://doi.org/10.1145/3351095.3372873

Royal Institute of British Architects. (2020). RIBA Plan of Work 2020. London: RIBA.

Royal Institute of British Architects. (2023). RIBA Health and Safety Guide (2nd ed.). London: RIBA Publishing.

Royal Institute of British Architects. (2025). RIBA Artificial Intelligence Report 2025. London: RIBA.

Schön, D. A. (1983). The reflective practitioner: How professionals think in action. Basic Books.