

Article

From AI Adoption to AI Transformation: The AX-5R Framework for Socio-Technical Work System Redesign

Kwan Soo Shin ^{1,2,3,*}, In Seok Kang ^{3,4} and Munho Lee ³¹ PolymathMinds Lab, Asan 31457, Republic of Korea² Department of AI and Innovation Management, Hanil University and Presbyterian Theological Seminary, Wanju 55359, Republic of Korea³ Graduate School, aSSIST University (Seoul School of Integrated Sciences and Technologies), Seoul 03767, Republic of Korea; iskang@postech.ac.kr (I.S.K.); munho419@naver.com (M.L.)⁴ Department of Chemical Engineering, Pohang University of Science and Technology (POSTECH), Pohang 37673, Republic of Korea

* Correspondence: sshin@pmminds.ai

Highlights

Please indicate how your work links to systems science via your contributions to systems practice, theory, and or methodology.

- AX-5R contributes to socio-technical systems theory and practice by integrating readiness, redesign, role, risk, and return into a workflow-level control loop that converts fragmented generative AI use into governed, measurable work system transformation.
- Adjacent frameworks supply the anatomy, stages, and controls of organizational AI adoption; AX-5R adds the workflow-level execution sequence that links them into one governed cycle.

What are the main findings and/or the implications of the main findings?

- In a blinded artifact-level ablation probe, the full framework outscored a sham five-part control and its component ablations, significant under two-sided Holm-corrected testing, and the machine scores converged with an independent human-expert-rating check ($r = 0.78$).
- Managers can move from AI pilots to governed scaling using the framework's implementation artifacts, configuration archetypes, and six testable propositions.

Abstract

Generative artificial intelligence (AI) has diffused rapidly, yet adoption has not reliably progressed to AI transformation (AX). Firms grant tool access but fail to redesign workflows, clarify accountability, govern risks, or measure value. The gap is a socio-technical systems problem, not a productivity problem: AI tools are inserted into existing routines without redesigning task interdependencies, decision rights, oversight loops, or performance feedback. This paper develops AX-5R, a socio-technical systems architecture that converts fragmented AI use into accountable, governable, and measurable work systems. Synthesizing seven literature streams, it maps failure modes to five interdependent design functions: readiness, redesign, role, risk, and return. AX-5R treats transformation as joint optimization of technical and social subsystems requiring configurational alignment. A supplementary ablation probe generated 252 artifacts across three workflows and seven prompt arms from two language models, scored by blinded cross-provider judges; the full frame outscored a sham five-part control and its four artifact-relevant ablations,

Academic Editors: Mathias Fonkam and Narasimha Rao Vajjhala

Received: 16 June 2026

Revised: 25 July 2026

Accepted: 26 July 2026

Published: 1 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

significant under two-sided Holm-corrected testing, with an independent human-expert-rating check. With a failure-mode derivation matrix, implementation artifacts, and six testable propositions, the framework specifies a minimum architecture in which readiness sets boundaries, redesign restructures tasks, role assigns accountability, risk establishes control, and return supplies learning feedback.

Keywords: AI transformation; generative AI; socio-technical systems; work system redesign; human–AI collaboration; AI governance; organizational transformation; AX-5R

1. Introduction

Generative AI has diffused rapidly into contemporary management practice. Large language models and related systems are now used for writing, summarization, coding, document search, customer response drafting, ideation, translation, analytics support, and internal knowledge work [1–4]. This diffusion has lowered the entry barrier to organizational AI use. Employees can now experiment with AI without waiting for a centralized enterprise system, a full data science team, or a large-scale IT transformation project.

However, this ease of adoption creates a paradox. The more easily AI tools are adopted, the easier it becomes for organizations to mistake usage for transformation. AI adoption means that individuals, teams, or departments use AI tools for particular tasks. AI transformation (AX), by contrast, begins when AI changes how work is designed, decisions are made, responsibilities are allocated, risks are governed, and performance is measured. A firm may have high AI usage but low AX performance. Employees may save time locally, yet the organization may not know which workflows changed, which risks increased, which human responsibilities were reallocated, or which outcomes improved.

This distinction matters because the value of AI depends less on isolated tool use than on organizational redesign. Digital transformation research shows that technology produces value when it changes processes, capabilities, routines, and business models [5–8]. Information systems research similarly demonstrates that adoption is not equivalent to assimilation, routinization, or effective use [9–11]. AI intensifies this older implementation problem because it can participate directly in cognitive work. Unlike earlier digital tools that mainly stored, transmitted, or processed information, generative AI can draft arguments, summarize documents, classify cases, recommend options, and simulate expertise [1,2,12,13]. The managerial challenge is therefore not simply to provide access to AI tools, but to redesign the surrounding work system to align with AI's appropriate role, accountability boundaries, risk exposure, and feedback logic.

The focal problem, therefore, is how organizations can convert fragmented AI use into workflow- or portfolio-level AX. This paper argues that such conversion requires five linked implementation moves: assessing workflow-specific readiness, redesigning workflows, clarifying human–AI roles, governing risk through human oversight, and measuring return through balanced indicators. These five moves form the AX-5R framework: readiness, redesign, role, risk, and return.

This gap is not merely a productivity problem but a failure of socio-technical systems integration. AI tools are added to individual tasks, but the surrounding work system, including workflow structure, role allocation, risk controls, and performance feedback, remains unchanged. As a result, local efficiency does not necessarily propagate into organization-level capability. Treated this way, AX is realized through the disciplined reconfiguration of an interdependent socio-technical system in which technical and social subsystems are jointly optimized, rather than a tool layered onto unchanged routines. This system's framing explains why fragmented adoption stalls and why improving any single

component in isolation leaves the larger system untransformed. It also clarifies why responsible oversight and balanced measurement are not optional add-ons: they are the control and feedback functions without which a work system cannot stabilize, scale, or sustain AI-enabled change.

This article is designed as an integrative framework study. It does not propose a new algorithm, estimate a causal model, or report a company case study. Instead, it synthesizes dispersed research streams into a socio-technical systems architecture for AI-enabled work system transformation. This is appropriate because current AX practice is ahead of theory integration. Managers face urgent questions: Which AI use cases should be selected first? Which workflows are sufficiently ready for AI-enabled redesign? Which tasks can be supported by AI, and which require human approval? How should risk be controlled as AI exposure increases? How should success be measured beyond vague productivity claims? The existing literature offers useful fragments, but practical AX requires an integrated architecture that links readiness, workflow redesign, role allocation, risk governance, and return measurement.

The study is guided by four conceptual research questions and one supplementary artifact-probe question:

RQ1. What explains the gap between AI adoption and workflow- or portfolio-level AX?

RQ2. Which literature streams are most relevant for theorizing AX as an implementation problem?

RQ3. How can organizations redesign workflows and human–AI roles while maintaining oversight?

RQ4. What practical framework can guide managers from AI pilots to scalable AX?

RQ5. In a supplementary artifact-level ablation probe using blinded cross-provider scoring, is the full AX-5R frame associated with higher-rated implementation artifacts than its single-component ablations and a sham five-part control?

The paper's central novelty lies not in adding another list of AI adoption factors but in specifying the socio-technical systems architecture required to convert fragmented AI tool use into redesigned, accountable, governable, and measurable organizational work systems. Unlike adoption-performance studies, governance-only models, broad reviews of AI's effects, or domain-specific workflow studies, this paper develops a cross-organizational execution architecture derived from recurring implementation failure modes. The contribution is fourfold. First, the paper separates AI adoption from AX and argues that AX begins only when work, responsibility, risk control, and performance systems are redesigned together. Second, the paper synthesizes seven literature streams that are often discussed separately. Third, the paper adds a supplementary artifact-level ablation probe using blinded cross-provider LLM scoring to examine whether the full AX-5R frame is associated with higher-rated implementation artifacts than its single-component ablations and a sham five-part control. Fourth, the paper provides implementation artifacts, including an AX lineage map, an adoption–transformation gap model, a human–AI role and risk matrix, an AX-5R process model, use-case selection criteria, role allocation levels, risk controls, and performance indicators.

2. Theoretical Background and Literature Review

2.1. Structural Overview

The AX problem sits at the intersection of seven studies. Digital transformation underscores the need to link technology adoption to strategic and organizational change. Information systems adoption explains why initial use does not automatically become assimilation, routinization, or effective use. AI and management research explains the automation–augmentation tension, the reallocation of decision authority, and the emergence

of human–AI collaboration as an organization-design problem. Generative AI research explains why general-purpose language systems differ from earlier enterprise systems: they diffuse informally, participate in cognitive work, and generate fluent but fallible outputs. The workflow redesign literature explains how technologies create value only when tasks, routines, and interdependencies are redesigned. The AI governance literature explains why human oversight, auditability, and risk controls must be embedded into work systems rather than added as external policy layers. The performance management literature explains why AX performance must be measured using balanced indicators rather than relying solely on tool usage or productivity anecdotes.

These streams are not separate bodies of literature but partial views of a single socio-technical systems problem. AI creates organizational value only when the technical artifact, workflow structure, human role system, governance controls, and feedback mechanisms are redesigned together. The central gap is therefore not the absence of research on AI adoption, digital transformation, or AI governance. It is the absence of an integrated systems architecture that explains how fragmented AI use can be converted into redesigned, accountable, governable, and measurable work systems.

In Figure 1, AX is positioned as the missing high-redesign and high-governance cell in prior research. Both axes are defined operationally. A workflow scores high on the horizontal axis when its before-and-after process maps differ in task structure, not merely in tool presence, and scores high on the vertical axis when review authority, input rules, escalation paths, and outcome indicators are documented parts of the workflow rather than external policies. The lower-left cell captures tool adoption and acceptance; the lower-right cell captures process redesign without sufficient oversight; the upper-left cell captures governance without work system transformation; and the upper-right cell identifies AX-5R as the integrated work-system architecture that closes the workflow, role, risk, and return gaps together.

AX requires both work-system redesign and governance/measurement integration.




Figure 1. AX is the missing cell in prior research. The horizontal axis ranges from insertion of AI into unchanged routines to restructured task interdependencies with explicit human–AI role allocation. The vertical axis ranges from ungoverned or informal use to oversight, escalation, and

measurement embedded in the workflow. The shaded upper-right cell marks AX-5R as the integrated work-system architecture in which both axes are satisfied jointly.

2.2. Literature Lineage and Benchmarking Logic

The review is organized as a lineage rather than a strict chronology. The first generation explains technology acceptance, implementation politics, IS success, effective use, situated practice, and the structuring effects of technology [9–11,14–23]. Its core lesson is that acceptance and use, by themselves, do not create organizational value.

Beneath this lineage lies a socio-technical systems foundation. Socio-technical systems theory holds that organizational performance depends on the joint optimization of interacting technical and social subsystems rather than on technology alone [24,25], and work system theory extends this principle to the deliberate redesign of the workflows, roles, information, participants, and technologies through which work is accomplished [26]. This foundation is directly relevant to AI, whose value depends on how the surrounding work system is reconfigured. However, generative AI sharpens the classic socio-technical problem because it diffuses through decentralized access to tools, participates in cognitive work, produces probabilistic outputs, and blurs the boundary between assistance, recommendation, and decision-making influence. Recent systems research applies the same lens to generative AI, conceptualizing AI adoption in firms as an integrated socio-technical system in which governance, capabilities, and technology jointly shape outcomes [27]. Parallel systems studies extend this lens to generative-AI-driven transformation in specific work domains [28] and show that balanced social and technical configurations, rather than technology alone, drive digital-transformation outcomes [29]. The present paper, therefore, does not merely restate the joint-optimization principle or estimate how AI adoption relates to firm performance; it operationalizes joint optimization as AX-5R, an execution architecture that converts fragmented AI use into a redesigned, accountable, governable, and measurable work system. Section 2.12 develops this positioning framework by framework.

The second and third generations shift from use to transformation capability. Digital strategy, digital innovation, platform transformation, IT-enabled organizational transformation, change management, and growth theory locate technology value in strategic response, organizational renewal, leadership, and complementary assets [5–8,30–40]. Capability, knowledge, process, and sociomaterial research then explain why firms differ in their ability to convert the same technology into repeatable routines and measurable work system change [41–55].

The fourth and fifth generations place AI inside task reallocation, organizational design, human–AI collaboration, and governance. Task-based economics and productivity studies show that AI effects vary by task, expertise, workflow integration, and complementary organizational capital [12,13,56–65]. Organizational AI research adds automation–augmentation tensions, conjoined agency, algorithmic control, hybrid intelligence, process automation, and meaningful-work risks [66–92]. Foundation models, risk, human factors, documentation, auditing, standards, and policy sources set the boundary condition: fluent AI output cannot be equated with reliable organizational action [1–4,93–112].

2.3. Digital Transformation: Technology Adoption Is Not Transformation

Digital transformation research treats technology value as an organizational change problem rather than a deployment problem. Digital technologies alter value creation, structures, processes, and business models only when firms combine digital capability, leadership, strategy, operating-model redesign, and dynamic capability formation [5–8,34,35,38,39]. Integrative reviews of digital transformation reach the same conclusion regarding the framework structure: technology-driven change succeeds through

interdependent managerial dimensions treated as a single configuration rather than through any single dimension, including in resource-constrained small and medium enterprises [113]. Applied to AX, the implication is direct: AI access without workflow redesign remains adoption; redesign without risk control remains fragile; and measurement based only on usage mistakes for value remains fragile. This stream settles that transformation is an organizational change problem, not a deployment problem. What it leaves unresolved is the task level: which workflow is redesigned, who verifies AI output inside that workflow, and what evidence justifies scaling all remain outside its unit of analysis [5–8,113]. AX-5R adopts the stream’s change thesis and supplies the missing workflow-level execution logic through Redesign and Return.

2.4. Information Systems Adoption: From Use to Assimilation and Effective Use

Information systems adoption research explains why acceptance is necessary but insufficient. TAM and UTAUT explain intention and use, while IS success, effective use, implementation, and practice-lens research distinguish adoption from assimilation, routinization, and net benefits [9–11,16,18,20,22]. This distinction is especially important for generative AI because informal experimentation can spread before policies, workflow maps, review standards, or performance measures are in place. AX therefore asks not only whether AI is used, but whether use becomes a redesigned and governable routine. This stream establishes the adoption–assimilation distinction on which the paper’s gap claim rests. It does not, however, explain how use converts into assimilation when the technology is fallible, decentralized, and cognitive, since its models predate that combination [9–11,16]. AX-5R treats that conversion mechanism as precisely the missing construct and names its components.

2.5. AI and Management: Automation, Augmentation, and Human–AI Collaboration

AI and management research frames AX as an automation, augmentation, and organization-design problem. AI can substitute for tasks, complement human capabilities, alter decision-making structures, generate joint agency, create data-network effects, and change competitive advantage [76,82,83,86,88]. Agentic IS, AI management, hybrid intelligence, organization design, and AI teammate research further show that delegation, autonomy, learning, inscrutability, and complementarity must be designed rather than assumed [68,70,74,81,84,87]. Trust and reliance studies add the human boundary: employees may overtrust fluent output or reject algorithms after errors, so AX requires explicit autonomy levels, review duties, escalation rules, and approval points [96,100,102,109]. The settled insight here is that human–AI configurations must be designed rather than assumed. What remains open is the implementation question: the stream specifies which tensions exist (automation versus augmentation, reliance versus rejection) but not the artifacts through which a specific workflow resolves them [68,70,83,88]. AX-5R answers with the Role component and its allocation matrix.

2.6. Generative AI and Foundation Models: A New Adoption Context

Foundation models make AX urgent because they are general-purpose, flexible, and accessible through natural-language interfaces [1,2]. Their broad capabilities support writing, summarization, coding, classification, and reasoning-adjacent work, but also create risks of plausible errors, fabricated sources, sensitive data exposure, and unreviewed AI-generated material [3,4,95,107,112]. Adoption-side research continues to accumulate at the level of use intention, including mixed-methods work on the determinants of generative AI use intention in domain-specific settings such as construction engineering [114]; such studies sharpen the explanation of why individual uptake spreads quickly, and equally sharpen the gap this paper addresses, because use intention is an individual-level

antecedent of adoption, not evidence of work system transformation. Generative AI therefore makes adoption more bottom-up and less centralized, increasing the need for rules about input, output, review, escalation, and measurement. This literature establishes the technology-side boundary condition: fluent output with plausible error, diffusing bottom-up through natural-language interfaces [1–4]. It does not provide an organizational response architecture, which is not its job. AX-5R takes the risk profile established here as a design input for Readiness and Risk.

2.7. Workflow Redesign and Technology-in-Practice

Workflow redesign research provides the mechanism for transformation. Process innovation and reengineering argue that technology value emerges when processes are redesigned rather than superficially digitized. At the same time, practice, affordance, and algorithm-at-work research show that technology effects depend on routines, identities, enactment, and control relations [19,20,42,46,48,75,80]. In AX, redesign decomposes work into retrieval, summarization, drafting, comparison, recommendation, approval, execution, and documentation, then assigns AI only to suitable subtasks. Robotic process automation research reinforces the same point: automation still requires decisions about which steps are scripted, monitored, escalated, or redesigned [90]. The redesign stream contributes the paper's core mechanism: value emerges from process restructuring, not tool insertion [19,20,42,46]. It is an open question for generative AI: which subtasks may be delegated to a probabilistic system and under what review? AX-5R resolves this by pairing Redesign with Role and Risk rather than treating redesign as a stand-alone move.

2.8. AI Governance, Human Oversight, and Risk Management

AI governance research shows that scaling increases exposure to risks of inaccuracy, bias, privacy, security, compliance, intellectual property, and accountability. NIST, the EU AI Act, and ISO/IEC 42001 [101] translate these risks into management-system expectations, while datasheets, model cards, audit frameworks, and accountability-gap research show why AI use needs documentation, traceability, and review artifacts [97,99,101,105–107,110]. Systems research similarly treats the governance of interconnected socio-technical systems as a design problem inside the system architecture rather than an external policy layer [115]. HR and meaningful-work research also add that governance must protect fairness, autonomy, employee reactions, and work quality [69,89]. In AX, governance is therefore proportional: low-risk summarization may need light-source checks, while legal, financial, employment, medical, safety, or regulated decisions require strict accountability or exclusion from automation. Governance research clarifies the control obligations and the documentation forms for them [97,101,105–107]. Left unspecified is the coupling problem: how control attaches to value-creating redesign instead of running as a parallel compliance program. AX-5R embeds proportional control within the workflow cycle, a design choice that distinguishes it from stand-alone governance frameworks (Section 2.12).

2.9. Performance Management and Value Realization

Performance management research prevents AX from becoming a usage narrative. Generative-AI productivity effects vary by task, expertise, review burden, and workflow integration, so time savings alone are insufficient [12,13,63]. Balanced-scorecard, IS success, and digital-transformation research point toward broader measurement: financial and operational efficiency, output quality, user satisfaction, customer impact, risk control, capability development, and strategic alignment [6,7,10,11,47]. AX therefore requires baselines, pilot outcomes, and learning indicators that tell managers whether to scale, revise, or stop a use case. This stream establishes that usage and time savings are not value evidence [10–13,47,63]. What it leaves open is which indicators can be used to adjudicate

a scaling decision for a specific redesigned workflow. AX-5R's Return component converts the balanced-measurement principle into pilot-level baselines and stop, revise, or scale rules.

2.10. Closing the Systems Gap: AX as Integrated Execution

Together, these studies explain why AI use must be converted into redesigned, governed, and measured work. Yet they remain fragmented across adoption, transformation, workflow, governance, and performance traditions. What remains missing is a systems-level implementation architecture that specifies how organizations move from fragmented tool use to accountable, governable, and measurable work system transformation. AX-5R is proposed to fill that gap. Stated critically rather than descriptively, the seven streams jointly explain why adoption fails to progress to transformation, but each stops one level short of execution. The strategy stream lacks the workflow unit; the workflow stream lacks the accountability and control logic; the governance stream lacks the value-creation coupling; and the measurement stream lacks the redesign object it is meant to measure. AX-5R is the configuration that simultaneously closes these four gaps, which is why it is derived from failure modes (Section 3.2) rather than assembled by juxtaposition.

In Figure 2, AX is presented as a five-generation lineage rather than a simple convergence diagram. The figure shows how IS adoption, digital transformation, work system redesign, AI management, and generative-AI governance accumulate toward the paper's focal claim: AX is realized through an integrated socio-technical systems architecture, AX-5R, rather than being a synonym for adoption or a stand-alone governance program. These generations are overlapping conceptual streams organized by their dominant contribution to AX, not strict chronological periods.

2.11. Construct Boundary: AX and Neighboring Concepts

Because AX is an emerging managerial term, it must be distinguished from neighboring concepts. Here, AX is not used as a synonym for AI adoption, digital transformation, AI governance, AI maturity, or responsible AI. These concepts overlap, but each emphasizes a different managerial problem. AI adoption focuses on whether AI tools are used. Digital transformation focuses on broader technology-enabled organizational change. AI governance focuses on risk, accountability, and control. AI maturity focuses on diagnostic assessment. Responsible AI focuses on ethical and social acceptability. AX, as used in this paper, focuses on the execution problem of converting AI tool use into redesigned work systems that have clear human–AI roles, proportionate oversight, and measured value.

This boundary is important for theoretical and practical reasons. If AX is treated only as adoption, organizations may measure usage and training while ignoring whether work has changed. If AX is treated only as governance, organizations may control risk while failing to redesign value-creating workflows. If AX is treated only as digital transformation, the distinctive challenges of generative AI (fluent but fallible output, shadow AI use, prompt-level data exposure, and ambiguous accountability) may disappear inside a broad transformation vocabulary. AX therefore occupies a middle position: it translates the strategic logic of digital transformation and the control logic of AI governance into an operational model for AI-enabled work redesign.




Figure 2. Five-Generation Literature Lineage for AX.

The construct boundary is summarized below.

In Table 1, AX is distinguished from neighboring concepts by shifting the unit of analysis from tool use, maturity assessment, governance principles, or general digital transformation to redesigned work systems. The table clarifies that AX is not a synonym for AI adoption; it is the realized work-system state that an organization's AX implementation capability produces by aligning workflows, roles, risks, and returns around AI-enabled work.

Table 1. Construct boundary of AX and neighboring concepts.

Concept	Primary Focus	Main Unit of Analysis	Typical Limitation	AX Distinction
AI adoption	Tool acceptance and use	Individual, team, or department	May stop at usage or experimentation	AX requires redesign of workflow, role, risk, and return systems
Digital transformation	Digital technology-enabled organizational change	Strategy, organization, business model	Often too broad to specify AI-specific roles and oversight problems	AX narrows transformation to AI-enabled cognitive and operational work
AI readiness	Preparedness for AI implementation	Data, capability, leadership, infrastructure	Often diagnostic rather than execution-oriented	AX uses readiness as the first step in an implementation cycle
AI maturity	Stage of AI capability development	Organization or function	May become a scoring model without workflow redesign	AX links maturity to concrete redesigned work and performance evidence
AI governance	Risk, compliance, accountability, and control	Policy, process, oversight structure	May emphasize control more than value creation	AX embeds governance inside workflow redesign and scaling
Responsible AI	Ethical, fair, transparent, and accountable AI	Principles, models, development practices	May remain principle-oriented or technical	AX translates responsibility into managerial artifacts and review routines
AI capability	Ability to deploy and use AI resources	Firm capability and resources	May under-specify role allocation and governance at the work level	AX treats capability as repeated work system execution across readiness, redesign, role allocation, risk control, and return feedback
Socio-technical systems theory	Joint optimization of social and technical subsystems	Work system or organization	Provides a general principle but not a generative-AI-specific implementation sequence	AX-5R operationalizes joint optimization for decentralized, cognitive, and fallible AI use through readiness, redesign, role, risk, and return
Work system theory	Workflows, roles, information, participants, and technologies	Work system	Broad work system lens without AI-specific accountability, oversight, and return logic	AX-5R specifies an AI-specific execution architecture for role allocation, risk control, and feedback-based scaling

The boundary table also defines the paper’s theoretical claim, which rests on a hierarchy of four related but distinct levels. AX-5R is a socio-technical work-system architecture: it specifies a workflow-level readiness gate and four transformation functions (redesign, role, risk, and return) through which fragmented AI use is converted into governed, accountable, and measurable work. AX implementation capability is the repeatable organizational ability to enact, adapt, and learn from this architecture across a portfolio of workflows. AX denotes the realized state of a workflow once minimum readiness conditions are met and the four transformation functions are enacted. AX performance refers to the measurable consequences of that realized state, and is not synonymous with AX

itself. AX-5R is therefore an architecture, not a capability, and the five Rs are its design functions rather than five parallel capabilities.

The construct can now be stated formally. AX is defined at the level of the work system, that is, an individual workflow or a portfolio of workflows, not at the level of the firm, the tool, or the individual user. Four conditions are individually necessary and jointly sufficient for a workflow to count as transformed rather than merely exposed to AI: (a) the task structure has been redesigned around AI-suitable subtasks rather than left unchanged (Redesign); (b) an accountable human role with explicit review or approval authority is allocated over AI-touched outputs (Role); (c) oversight and escalation controls proportionate to the workflow's stakes and data sensitivity are in place (Risk); and (d) outcomes are measured against a pre-AI baseline on more than usage or time saved (Return). Readiness is deliberately excluded from the definition: it is the entry condition that makes satisfying the four definitional conditions feasible, not a component of the transformed state itself, which is consistent with its behavior in the ablation probe (Section 3.3). Each condition has a directly observable indicator, namely the implementation artifacts of Section 5: a before-and-after workflow map for (a), a completed role allocation matrix with named approval points for (b), documented input rules, source checks, audit logs, and escalation rules for (c), and a baseline and post-pilot indicator dashboard for (d). The indicator also provides each condition's threshold: a condition is satisfied when its artifact exists and is enacted in the live workflow, not merely documented. A workflow satisfying none or only some of the four conditions instantiates adoption, partial integration, or governance without transformation, as located in Figure 1.

Readiness is architecturally integral to AX-5R but ontologically antecedent to the realized AX state. It is the workflow-level entry and re-entry condition for the redesign, role, risk, and return transformation core: it determines whether a workflow is eligible to enter the transformation cycle, and it is reassessed when return feedback reveals new constraints, drift, or capability gaps. Readiness is therefore one of the five design functions of the architecture, but it is not one of the four conditions that define an achieved AX state. Because the artifact-level ablation probe (Section 3.3 and Appendix A) evaluates implementation artifacts after the AX-5R framing is supplied, it was not designed to test workflow eligibility directly; the non-significant readiness contrast is therefore consistent with, but does not confirm, this placement of readiness as an entry condition.

The seven literature streams that inform this synthesis are summarized below.

In Table 2, the seven literature streams are translated into implementation requirements. The table shows that each stream contributes one necessary part of AX: organizational change, effective use, human–AI role design, generative AI risk awareness, workflow redesign, governance, and balanced performance measurement.

Table 2. Seven literature streams as inputs to AX-5R derivation.

Stream	Core Insight	AX Implication
Digital transformation	Technology value requires organizational change	AI tool use is not AX
IS adoption	Use differs from assimilation and effective use	AI must become a redesigned routine
AI and management	Automation and augmentation coexist	Human–AI roles must be explicit
Generative AI	Flexible tools diffuse informally	Shadow AI and review risk must be managed
Workflow redesign	Process change creates technology value	AI use cases must become workflow maps
AI governance	Risk must be governed and documented	Oversight is part of implementation
Performance management	Value requires measurement	AX needs balanced KPIs

2.12. Positioning AX-5R Against Adjacent Frameworks: What Generative AI Changes

The construct boundary in Table 1 separates AX from neighboring concepts. A reviewer of any new framework, however, may still ask a sharper question: given that socio-technical systems theory, work system theory, digital transformation frameworks, AI maturity models, and AI governance frameworks already exist, what exactly does AX-5R add? This subsection answers that question directly, at the level of framework architecture rather than concept labels.

Work system theory is the closest neighbor. Alter's work system framework decomposes any work system into elements, including workflows, roles, information, participants, and technologies [26], and its granularity is deliberately close to AX-5R. The two frameworks nevertheless answer different questions. Work system theory provides the anatomy of a work system: it tells analysts what a work system consists of. AX-5R specifies the transformation control loop for one specific technology class: it tells managers in what order, under what entry conditions, with what accountability allocation, under what control loops, and against what feedback evidence a generative-AI-enabled redesign should proceed. Anatomy does not entail a control loop, and a control loop is exactly what the adoption–transformation gap demands.

Generative AI motivates this control loop through five design considerations that pre-generative frameworks did not need to encode. First, diffusion is decentralized and bottom-up: employees adopt tools through natural-language interfaces before any formal implementation project exists, so adoption routinely precedes governance rather than following it [1,2]. Second, output is fluent but probabilistic: plausible error, fabricated sources, and unstable answers make verification a designed task rather than an optional courtesy [3,4,95,107]. Third, the technology participates directly in cognitive work: drafting, summarizing, classifying, and recommending, so the accountability boundary runs inside tasks rather than between jobs [12,13,83,88]. Fourth, use is prompt-mediated: sensitive data can leave the organization one prompt at a time, which moves input governance from the IT perimeter into everyday work practice [97,107]. Fifth, capability drift: vendor-side model updates silently change system behavior, so a workflow validated once is not permanently validated, and re-validation must be looped into performance feedback [93]. Readiness, redesign, role, risk, and return are the five capabilities that respectively absorb these considerations; a framework lacking any one of them leaves the corresponding consideration unmanaged.

Table 3 positions AX-5R against each adjacent framework family on these terms.

Table 3. Positioning AX-5R against adjacent framework families.

Framework Family	Primary Purpose	Unit of Analysis	What It Specifies	What It Leaves Unspecified for Generative AI	Relation to AX-5R
Socio-technical systems theory [24,25]	Explain the joint optimization of social and technical subsystems	Organization or work system	The principle that technology and social structure must be redesigned together	The implementation sequence, artifacts, and control loops for decentralized, probabilistic, cognitive AI	AX-5R operationalizes the joint-optimization principle as an executable architecture

Work system theory [26]	Describe and analyze work systems	Work system	The anatomy: workflows, roles, information, participants, technologies	The transformation control loop: entry conditions, accountability allocation, proportionate oversight, and feedback for a fallible cognitive technology	AX-5R adds the generative-AI-specific execution sequence to the work system lens
Digital transformation frameworks [5–8,34,35,113]	Explain technology-enabled organizational change	Strategy, organization, business model	Strategic drivers, capability building, and integrative change dimensions	Task-level role allocation, verification burden, prompt-level data exposure, and model-drift governance	AX-5R narrows transformation logic to AI-enabled work system execution
AI maturity models [116]	Assess stages of AI capability	Organization or function	Diagnostic stage descriptions and scoring dimensions	The mechanism that moves a firm between stages; maturity scores do not redesign any workflow	AX-5R consumes the maturity assessment as Readiness input and specifies the movement mechanism
AI governance frameworks [97,101,106,107,110]	Control AI risk and ensure accountability	Policy, management system, oversight structure	Risk taxonomies, documentation, audit, and management system requirements	Where value is created: which workflows to redesign, how roles change, and how return is measured	AX-5R embeds proportionate governance inside the redesign cycle instead of alongside it
AX-5R (this paper)	Convert fragmented AI use into a governed, measurable work system transformation	Workflow and workflow portfolio	Entry conditions, task restructuring, accountability allocation, control loops, and outcome feedback as one configuration	Field validation across industries remains a research agenda (Section 5.8)	Integrates the five families into one execution architecture

In Table 3, each adjacent framework family contributes a necessary layer, and none is displaced. The specific claim of AX-5R is integrative and executorial: it is the only layer in this stack that specifies a transformation control loop at the level of the individual workflow, with components that map onto recurring generative-AI failure modes (Section 3.2) and outputs that are implementation artifacts rather than principles, stages, or policies. This is also why five components are specified: each absorbs one of the five generative-AI design considerations above and corresponds to one recurring failure mode derived in Section 3.2, so removing a component leaves a failure mode without a designated owner, while adding a sixth would re-partition functions already assigned rather than absorb a new failure mode. This mapping is a design argument rather than a logical necessity, and the five components do not all operate at the same level: four are embedded in implementation artifacts, whereas Readiness is the organizational entry condition that the construct definition in Section 2.11 deliberately places before the transformed state. The ablation probe in Section 3.3 tests the non-redundancy of the artifact-embedded components at the artifact level, and its results align with this division: removing any artifact-embedded component was penalized, whereas removing Readiness was not.

3. Research Method

This study uses an integrative literature synthesis and framework design approach, supplemented by an artifact-level ablation discriminant probe. The objective is not to estimate a causal effect, conduct a systematic review of a narrowly bounded domain, or report a company case study. Rather, the objective is to synthesize dispersed research streams into a socio-technical systems architecture for AI-enabled work system transformation. This design is appropriate because AX is still an emerging construct and its relevant knowledge is distributed across digital transformation, information systems adoption, AI management, generative AI, workflow redesign, governance, and performance management.

3.1. Integrative Literature Synthesis

The synthesis protocol is summarized below.

Table 4 makes the review procedure traceable. The review was not designed as a systematic review or meta-analysis; it was a purposive, theory-building synthesis aimed at identifying recurring implementation failure modes and translating them into a socio-technical systems architecture. The final corpus contains 117 sources, including peer-reviewed conceptual and empirical studies, influential books and working papers, technical papers on the foundations of generative AI, and official standards or policy documents used as boundary-setting evidence. A source was included when it explained technology-enabled organizational transformation; distinguished adoption, use, assimilation, or effective use; analyzed AI's effect on work, decision-making, or organizational capability; addressed generative-AI risks, oversight, accountability, or auditability; or offered performance-measurement logic relevant to implementation. Sources were excluded when they focused solely on technical model architecture without organizational implications, offered speculative commentary without clear conceptual value, or lacked relevance to implementation, governance, work redesign, or performance measurement.

Table 4. Integrative review and framework synthesis protocol.

Element	Procedure	Output for AX Framework Development
Review purpose	Concept development and implementation framework design, rather than effect-size estimation	A construct boundary and implementation logic for AX
Database and discovery routes	Seed works, backward and forward snowballing, publisher metadata, Crossref/DOI verification, Google Scholar and Semantic Scholar discovery searches, standards, and policy databases	A cross-stream corpus rather than a single narrow review domain
Search strings and search families	"digital transformation", "information systems adoption", "effective use", "AI adoption", "generative AI work", "human–AI collaboration", "AI governance", "AI risk management", and "performance measurement"	Seven literature streams linked to AX implementation
Search window	Discovery searches and snowballing were conducted in June 2026, with no lower publication date bound, so that foundational works remained eligible; the corpus spans 1951 to 2026	A corpus combining foundational theory and current generative-AI evidence

Inclusion criteria	Sources explaining organizational transformation, assimilation/effective use, AI-enabled work, human–AI roles, governance, risk, or performance measurement	The literature relevant to moving from AI use to organization-level transformation
Exclusion criteria	Technical model papers without organizational implications, commentary without stable conceptual value, speculative claims without implementation relevance	Reduced risk of treating AI enthusiasm as evidence
Coding focus	Failure mode, unit of analysis, managerial requirement, practical artifact, boundary condition, and testable implication	Failure-mode-to-capability mapping
Coding procedure	The first author coded all sources against the six coding foci, which served as the coding instrument, consistent with an integrative rather than systematic review methodology	Failure-mode-to-capability mapping traceable to sources
Consistency check	The co-authors independently reviewed and confirmed the stream assignment and failure-mode mapping for all Table 5 entries; any disagreements were resolved by discussion until consensus	Verified mapping between streams, failure modes, and AX-5R components
Final synthesis output	117 references organized around seven streams and five recurring implementation requirements	AX-5R: Readiness, redesign, role, risk, and return

Three methodological clarifications address the transparency expectations for integrative reviews. First, the search window: discovery searches and snowballing were conducted in June 2026 with no lower publication date bound, so the corpus spans 1951 to 2026 and combines foundational theory with current generative-AI evidence. Second, coding: the review followed an integrative-synthesis methodology, in which the coding instrument is a structured coding focus rather than a systematic review codebook; the six coding foci in Table 4 served that function, and every framework-bearing claim in Table 5 is traceable to cited sources. Third, consistency: the co-authors independently reviewed and confirmed the stream assignment and failure-mode mapping, with any disagreements resolved by discussion until consensus. The review is therefore reported as what it is, a purposive theory-building synthesis, and the manuscript does not claim systematic-review completeness; the corresponding completeness limitation is stated in Section 3.4.

The synthesis proceeded in four stages. First, the literature was grouped into seven streams. Second, recurring implementation failure modes were extracted from those streams. Third, the failure modes were translated into managerial requirements. Fourth, the requirements were integrated into five interdependent AX-5R components: readiness, redesign, role, risk, and return.

3.2. Framework Derivation Logic

The AX-5R derivation logic is summarized below.

Table 5. Derivation of AX-5R from recurring AI implementation failure modes.

Recurring Failure Mode	Literature Root	AX-5R Response	Observable Implementation Artifact
Tool access without organizational preparedness	IS adoption, assimilation, absorptive capacity	Readiness	Readiness diagnosis and workflow suitability screen

AI use without process-level change	Digital transformation, workflow redesign, and effective use	Redesign	Before/after workflow map and re-designed task sequence
AI output without an accountable human role	Human–AI collaboration, automation reliance, decision authority	Role	Human–AI role allocation matrix and approval protocol
Experimentation without controlled scaling	AI governance, risk management, standards, and auditability	Risk	Input rules, source checks, audit logs, and escalation rules
Productivity claims without balanced evidence	IS success, balanced scorecard, value realization	Return	Baseline/post-pilot dashboard and learning review

In Table 5, the five AX-5R components are traced to recurring implementation failure modes in the literature. The table shows that AX-5R is not introduced as a mnemonic list but derived as a problem-to-requirement architecture. Readiness responds to tool access without organizational preparedness, Redesign to AI use without process-level change, Role to AI output without accountable human responsibility, Risk to experimentation without controlled scaling, and Return to productivity claims without balanced evidence.

This synthesis used problem-to-requirement logic. Readiness identifies whether the organization and workflow can absorb AI-enabled change. Redesign converts individual AI use into process-level change. Role clarifies human–AI responsibility. Risk embeds proportional oversight. Return evaluates whether the redesigned work system creates measurable value and transferable learning. In systems terms, these components correspond to boundary and precondition diagnosis, task-interdependence restructuring, agency and accountability allocation, control and escalation loops, and outcome feedback.

3.3. Supplementary Ablation Discriminant Probe

As a supplementary structural test of the framework, the study ran an ablation discriminant probe without using human subjects. The probe is an AI-augmented structural scenario test in which controlled prompt conditions generate comparable implementation artifacts. It tests whether the full AX-5R frame produces distinguishable implementation artifacts under matched task conditions, not whether AX-5R improves organizational performance.

Two cross-provider language models, gpt-4o-mini and claude-haiku-4-5-20251001 (hereafter Claude Haiku 4.5), generated 252 implementation artifacts across three workflows and seven prompt arms. The three workflows were customer-service response, internal management reporting, and HR document screening. The seven arms were the full AX-5R frame, five single-component ablations that each removed one of readiness, redesign, role, risk, and return, and a sham five-part control consisting of a plausible but transformation-irrelevant five-step method.

To reduce same-provider circularity and label-conformity bias, each artifact was scored independently by both provider models. For the primary analysis, each artifact's score was taken from the judge belonging to the provider opposite to the generator: OpenAI-generated artifacts were scored by the Anthropic model, and Anthropic-generated artifacts were scored by the OpenAI model. Judges were blind to the prompt arm and to any framework name. No second artifact was placed in the judge's context, eliminating pairwise position-order bias.

Each artifact was scored on five fixed implementation-quality criteria (detailed in Appendix A) and on a single holistic implementation-quality score from 0 to 10, which served as the primary outcome. The criteria captured generic implementation-quality

dimensions derived from the literature synthesis rather than explicit conformity to AX-5R labels.

Because the two generators differed in baseline verbosity and scoring level, holistic scores were standardized within each generator before pooling, thereby isolating framing effects from generator-level score shifts. All contrasts were also estimated within each generator separately. For each arm, the full-versus-arm difference was summarized using Cohen's *d* with 95% confidence intervals, and the Mann–Whitney test was used to assess the difference. Because the probe tested a directional discriminant expectation specified in advance (the full framework should receive higher implementation-quality ratings than ablated or sham arms), one-sided tests were pre-specified; however, all results for these contrasts are reported two-sided in Table 6, which is the more conservative convention, and the substantive conclusions are identical under either convention. The five pre-specified contrasts (the sham control and the four artifact-relevant ablations) were corrected for multiple comparisons with the Holm–Bonferroni procedure; all five remained significant after correction under both the one-sided and the two-sided convention (largest Holm-adjusted two-sided $p = 0.042$). The Readiness contrast is reported separately as a scope-limited contrast outside the pre-specified set for the reason given below. Per-arm descriptive statistics (means, standard deviations, 95% confidence intervals, and group sizes, both raw within each generator and pooled after within-generator standardization) are reported in Appendix A Table A2.

The cross-provider score correlation between the two providers was moderate (Pearson $r = 0.66$, $n = 251$). As an additional check on the machine-judged scores, three human domain experts, none of whom is an author, independently rated a stratified subsample of 56 artifacts (8 per arm, balanced across generators and near-balanced across the three workflows) under the same anchored rubric, blind to arm labels and to the framework's identity; the mean expert rating converged with the cross-provider machine scores ($r = 0.78$), and full inter-rater reliability and convergence results are reported in Appendix A.8.

The probe results are summarized below.

Table 6. Ablation discriminant probe results (full AX-5R versus each arm; pooled within-generator standardized scores; $n = 36$ per arm).

Contrast (Full vs. Arm)	Cohen's <i>d</i>	95% CI for <i>d</i>	<i>p</i> (Two-Sided)	Holm-Adjusted <i>p</i>	Reading
Sham five-part control	2.46	[1.84, 3.07]	<0.001	<0.001	Full framework scores far above a five-part scaffold of equal form, but no transformational content
Redesign removed	1.01	[0.52, 1.50]	<0.001	<0.001	Largest single-component drop; consistent with redesign as the mediating mechanism (Proposition 2)
Risk removed	0.70	[0.23, 1.18]	0.004	0.012	Risk ablation receives lower ratings
Role removed	0.43	[−0.04, 0.89]	0.021	0.042	Role ablation receives lower ratings

Return removed	0.49	[0.02, 0.96]	0.038	0.042	Return ablation receives lower ratings
Readiness removed (scope-limited)	−0.28	[−0.74, 0.19]	0.124	n/a	No artifact-level penalty; consistent with Readiness as an organizational precondition rather than an artifact-embedded feature

Note: n/a indicates that the readiness contrast lay outside the pre-specified family of five contrasts and was therefore not entered into the Holm–Bonferroni correction. p -values are from Mann–Whitney tests on the pooled standardized scores, while the confidence intervals are normal-approximation intervals for Cohen’s d . Because the test is rank-based and the interval is parametric, an interval that crosses zero can accompany a significant rank-based p , as in the role contrast. On the rank scale matched to the test, the rank-biserial correlations (bootstrap 95% confidence intervals, 10,000 resamples) are 0.88 [0.78, 0.95] for the sham contrast, 0.58 [0.39, 0.75] for redesign, 0.37 [0.15, 0.59] for risk, 0.30 [0.06, 0.53] for role, and 0.27 [0.02, 0.50] for return; all four ablation intervals exclude zero, while the readiness interval, outside the pre-specified set, does not (−0.20 [−0.44, 0.04]).

In Table 6, the full AX-5R frame received higher ratings than a sham five-part control by a large margin (Cohen’s $d = 2.46$, $p < 0.001$), which suggests that the observed difference is more consistent with the framework’s substantive content than with merely presenting five labeled sections. Removing any of the four artifact-embedded components, redesign, role, risk, or return, received lower blind-scored implementation-quality ratings (all two-sided $p < 0.05$), and redesign removal showed the largest rating drop ($d = 1.01$), consistent with workflow redesign as the mediating mechanism in Proposition 2. The artifact-embedded components were therefore not cost-free to remove, consistent with the configurational claim in Proposition 6. In contrast, readiness behaved as an organizational entry condition rather than an artifact-quality feature: removing readiness did not lower artifact quality, and the non-significant point estimate was, if anything, slightly negative ($d = -0.28$, two-sided $p = 0.124$, meaning the readiness-removed arm scored nominally higher), consistent with its function as a pre-implementation organizational condition rather than an artifact-embedded design feature, so the artifact-level probe is not expected to be equally sensitive to it. The direction of the artifact-relevant contrasts was positive within both generators, although the strength of individual contrasts varied by generator; the pooled within-generator standardized analysis therefore serves as the primary summary, with the per-generator contrasts reported as robustness diagnostics. These contrasts are based on generator-standardized holistic scores with opposite-provider scoring. All five pre-specified contrasts remained significant after Holm–Bonferroni correction under the conservative two-sided convention (largest adjusted $p = 0.042$), and the pre-specified one-sided analysis yields the same conclusions with smaller p -values. Full descriptive statistics for every arm appear in Appendix A Table A2.

3.4. Evidence Boundary and Validity Limits

The study’s evidence burden is conceptual coherence, the literature coverage, construct boundary clarity, failure-mode derivation, practical usability, worked application logic, a supplementary artifact-level probe, and testable propositions. AX-5R has not yet been field-validated against organizational performance outcomes. The workflow vignettes and the ablation discriminant probe provide artifact-level evidence rather than field- or organizational-level performance evidence. Later empirical work should examine

whether AX-5R predicts AX performance across firms, industries, tasks, and regulatory settings.

To state the evidential status of each claim type precisely: the construct boundary (Section 2.11) and the framework architecture (Section 4) rest on the literature synthesis; the non-redundancy of the artifact-embedded components rests on the ablation discriminant probe, which is artifact-level, model-generated, and model-judged evidence with an added human-expert-rating check (Appendix A.8); and the six propositions are testable claims whose organizational validation is the research agenda of Section 5.8, not a result of this paper. The probe was designed so that it could fail: a sham control with identical surface form provided the discriminant test, and one component (readiness) indeed showed no artifact-level penalty, an outcome the framework accommodates by assigning readiness to the organizational entry stage rather than to artifact-embedded design. A further interpretive limit follows from the design itself: four of the five scoring criteria are intentionally aligned with the artifact-embedded components, so lower scores under ablation partly reflect that definitional alignment rather than an independent validity test. The contrasts that carry inferential weight are therefore the sham control, which shows that a plausible but transformation-irrelevant five-part scaffold does not earn equivalent ratings, and the human expert layer (Appendix A.8), which checks that the machine-scored separation is not an artifact of model judging. No claim in the manuscript asserts field-validated effects on organizational performance.

Except for Table 6 and Appendix A Table A2, all figures and tables in the manuscript are author-developed conceptual and implementation artifacts derived from the literature synthesis. Table 6 is presented as a pre-specified, artifact-level check of prompt-structure fidelity and internal consistency, scored by independent cross-provider model judges, not as mechanism or validity evidence at the organizational level.

4. AX-5R Implementation Framework

The AX-5R framework theorizes AI transformation as an iterative socio-technical work-system architecture. AI tools create organizational value only when organizations assess readiness, redesign work, allocate roles, govern risk, and measure value and learning. These moves correspond to five recurring failure points: access without preparedness, use without process change, output without accountability, experimentation without controlled scaling, and activity without value evidence. AX-5R is therefore not only a process model but a socio-technical work-system architecture for converting fragmented AI use into redesigned, accountable, governable, and measurable work systems. The overall sequence is as follows:

AI tool adoption → workflow-specific readiness → workflow redesign → human–AI role clarity → proportional risk control → balanced return measurement → scalable AX performance.

The order is sequential but not strictly linear. Role ambiguity, risk exposure, performance evidence, and implementation learning can feed back into readiness diagnosis and workflow redesign.

In Figure 3, AX-5R is shown as a readiness-gated transformation cycle rather than a set of five parallel components. Readiness sits outside the internal cycle as an entry and re-entry gate that determines whether a workflow is eligible to proceed. Within the gate, redesign, role, risk, and return form the transformation core: redesign restructures the workflow, role allocates human and AI authority, risk sets proportional oversight, and return supplies performance and learning feedback that re-enters readiness and redesign. The figure also locates the six propositions: Proposition 1 governs entry at the gate; Proposition 2 is the mediation path from adoption through redesign to AX performance; Propositions 3 and 4 are moderation relationships acting on that path; Proposition 5 is the

return feedback loop; and Proposition 6 is the configurational fit that binds the four transformation functions.




Figure 3. AX-5R socio-technical architecture: readiness gate, transformation core, and proposition topology. Solid arrows show the forward execution sequence; dashed arrows mark the mediation path to AX performance (Proposition 2); and the curved arrow is the return feedback loop (Proposition 5) re-entering the readiness gate.

4.1. Readiness

Readiness is the organization’s workflow-specific preparedness to use AI responsibly and effectively. It includes data, process, employee, leadership, measurement, and governance readiness. The key question is not whether AI access exists, but whether a workflow has the documents, review expertise, baseline measures, data rules, and policy maturity needed to absorb AI use into routines. Readiness prevents premature scaling by identifying workflows where data are unavailable, review standards are weak, accountability is unclear, or baseline outcomes cannot be measured. In systems terms, readiness defines the boundary and preconditions of AI-enabled transformation.

4.2. Redesign

Redesign means changing the workflow rather than adding AI to an unchanged process. It maps the current and AI-supported process, specifies what AI produces, what evidence it uses, who reviews it, what criteria apply, and where approval occurs. By decomposing work into retrieval, summarization, drafting, classification, comparison, recommendation, approval, execution, and documentation, redesign converts a vague use case into a work architecture. It is the central mediating mechanism through which AI adoption progresses to AX rather than mere tool use. In systems terms, redesign restructures task interdependencies.

4.3. Role

Role refers to the allocation of task authority, review responsibility, and accountability between humans and AI systems. AI may retrieve, summarize, draft, classify, recommend, or execute, but higher autonomy requires stronger oversight. AX separates task authority from accountability authority: AI may draft or recommend, while humans

remain accountable for review, exception handling, approval, and stakeholder consequences. Role clarity moderates redesign effectiveness by reducing ambiguity about when to trust, verify, override, escalate, or document AI output. In systems terms, role allocates agency and accountability.

4.4. Risk

Risk refers to the control architecture governing inputs, outputs, processes, and incidents as AI exposure scales. Input controls specify permissible data, output controls specify review requirements, process controls specify logging, and incident controls specify escalation. Risk control is proportional: internal summaries may need light-source checks, customer-facing drafts may need stronger review, and high-impact legal, financial, employment, medical, safety, or regulated decisions may require strict human accountability or exclusion from automation. This embeds governance into the work routine rather than leaving it as a policy document. In systems terms, risk establishes control and escalation loops.

4.5. Return

Return means measuring whether the redesigned workflow creates value and generates transferable learning. It is broader than financial ROI and includes efficiency, quality, human capability and burden, risk reduction, customer or stakeholder impact, and organizational learning. Baseline and post-redesign indicators allow managers to scale, revise, or stop a use case. A pilot that saves time but increases errors needs redesign; a pilot that creates local productivity but no transferable learning remains an act of adoption rather than transformation. In systems terms, return supplies performance feedback and learning.

4.6. Interdependence of the Five Components

The five components are mutually reinforcing and systemically interdependent. Readiness without redesign remains diagnosis; redesign without role clarity creates accountability gaps; role clarity without risk control may fail in high-stakes settings; risk control without return measurement may become bureaucracy; and return measurement without readiness and redesign may capture local tool effects rather than transformation. AX-5R therefore treats scalable AX performance as configurational alignment across the four transformation functions once minimum readiness conditions are met, not as excellence in any single component.

In systems terms, each component performs a distinct function. Readiness defines system boundaries and preconditions; redesign restructures task interdependencies; role allocates agency and accountability; risk establishes control and escalation loops; and return supplies performance feedback and learning. AX-5R is therefore a functional socio-technical architecture for AI-enabled work system transformation rather than a checklist of adoption factors.

Interdependence does not mean harmony. The five components generate managed tensions, and naming them is part of the framework's realism. The sharpest is between risk and redesign: control slows redesign agility, and redesign pressure erodes control discipline. AX-5R does not dissolve this tension; it prices it through proportionality, so that low-stakes workflows buy agility with light controls while high-stakes workflows deliberately trade speed for accountability (Section 6.3). A second tension runs between return and experimentation: premature measurement can kill exploratory pilots, while unmeasured pilots never earn the right to scale; the baseline-then-pilot sequencing in Phase 2 exists to stage this trade-off rather than deny it. A third tension links readiness and momentum: readiness gates protect the organization, but can become a veto point

that starves transformation of early wins. These tensions are not design flaws. They are the reason configurational alignment (Proposition 6) is the framework's success criterion: an aligned configuration is one in which each tension has been explicitly priced for that workflow's stakes, rather than resolved by default in favor of speed or of control.

5. Implementation Artifacts and Propositions

This section translates the AX-5R conceptual systems architecture into implementation artifacts and testable propositions. The artifacts are not presented as field-validated instruments or universal managerial checklists. Rather, they operationalize the systems functions specified in Section 4: boundary and precondition diagnosis, task-interdependence restructuring, agency and accountability allocation, control and escalation loops, and outcome feedback. Their purpose is to show how the AX-5R architecture can be applied to use-case selection, human–AI role allocation, risk control, performance measurement, staged scaling, and future empirical research.

5.1. AI Adoption Versus AX

In Table 7, AI adoption and AI transformation are separated by their respective units of analysis and managerial logics. Adoption remains centered on tool access and local use, whereas AX requires redesigned workflows, explicit oversight, and a performance system that can distinguish activity from organizational value.

Table 7. AI adoption versus AI transformation.

Dimension	AI Adoption	AI Transformation
Primary unit	Tool, task, or user	Workflow, role, governance, and performance system
Change type	Local use	Organizational redesign
Managerial focus	Access and training	Execution and scaling
Risk treatment	Informal discretion	Explicit oversight
Performance logic	Productivity anecdotes	Balanced indicators
Typical failure	Fragmented use	Redesign-governance-performance misalignment

5.2. Use-Case Selection

In Table 8, AX use-case selection is framed as a managerial screening problem. The table recommends beginning with frequent, document-based, reviewable, measurable, and lower-risk workflows because these characteristics enable comparison of baseline and post-redesign performance while controlling implementation risk.

Table 8. AX use-case selection criteria.

Criterion	Good First AX Candidate	Poor First AX Candidate
Frequency	High-volume repeated work	Rare exceptional work
Knowledge base	Source documents available	Tacit undocumented knowledge
Reviewability	Output can be checked	Output cannot be verified easily
Risk	Low to moderate consequence	High legal, safety, employment, or financial consequences without mature controls
Workflow clarity	Process can be mapped	Process is ambiguous
Data sensitivity	Low sensitivity	Confidential or regulated data
Measurement	Baseline available	No measurable outcome

5.3. Human–AI Role Allocation

In Table 9, human–AI role allocation is organized by increasing AI task authority. The table makes explicit that the human role changes as AI moves from retrieval and


summarization to drafting, recommendation, and bounded autonomous action, but accountability does not disappear; it must be reassigned, reviewed, or restricted.

Table 9. Human–AI Role Allocation Matrix.

Level	AI Role	Human Role	Suitable Use
L0	No AI	Full human execution	High-stakes decisions
L1	Retrieval	Source selection and verification	Policy/document search
L2	Summarization	Completeness and accuracy review	Meeting or report summary
L3	Drafting	Edit, verify, approve	Email/report/customer response draft
L4	Recommendation	Evaluate alternatives and decide	Vendor or option comparison
L5	Bounded autonomous action	Monitor, audit, and retain override authority	Pre-approved low-risk routine action only

In Figure 4, human–AI role allocation is linked to proportional risk governance. AI exposure is defined as the combined level of AI autonomy and task consequentiality, while the vertical axis represents the strength of human accountability and oversight. Low-exposure assistive work, such as retrieval or summarization, can be governed through light but explicit review. As use cases move toward recommendation, decision support, agentic execution, or high-impact tasks, they must move upward toward stronger accountability, approval, auditability, and escalation. The scalable AX zone is reached only when increasing AI exposure is matched by proportionate human accountability and oversight. Otherwise, high-exposure AI use becomes shadow AI: technically capable but poorly governed.

Scalable AX requires accountability and oversight to scale with AI exposure.



Diagnostic rule: move use cases up as they move right; otherwise high-exposure AI becomes unmanaged exposure.

Figure 4. AX Risk-Responsibility Diagnostic Map.

5.4. Risk Control Checklist

In Table 10, AI governance is operationalized as concrete controls rather than abstract principles. The table links each risk area to a control question and a practical artifact, allowing managers to embed oversight into everyday workflow design instead of treating governance as a separate compliance layer.

Table 10. AX risk controls.

Risk area	Control Question	Practical Artifact
Accuracy	Can the output be verified against sources?	Source-check rule
Privacy	Does the prompt contain personal or confidential data?	Input policy
Security	Is the tool approved for business use?	Approved-tool list
Bias	Could output disadvantage a stakeholder group?	Bias review checklist
Accountability	Who approves the final output?	Approval matrix
Auditability	Can AI use be traced if needed?	Use log
Incident response	What happens when AI output causes a problem?	Escalation rule
Learning	Are incidents and corrections fed back into workflow redesign?	Post-incident learning log
Provenance/IP	Are sources, ownership, and reuse permissions clear?	Source/provenance rule

5.5. Performance Indicators

In Table 11, AX performance is measured through a balanced set of indicators. The table prevents a narrow interpretation of return as time savings alone by adding indicators of quality, customer, human, risk, and learning that can reveal whether AI-supported work is genuinely improving or merely becoming faster. In systems terms, return indicators function as feedback signals that determine whether a redesigned workflow should be scaled, revised, or stopped.

Table 11. AX performance indicators.

Category	Indicator	Interpretation
Efficiency	Time saved, cycle time, throughput	Process speed and productivity
Quality	Error rate, correction rate, reviewer burden	Output reliability
Customer	Response speed, satisfaction, and resolution rate	External value
Human	Confidence, skill development, perceived role clarity, workload strain	Employee capability, role clarity, and sustainable workload
Risk	Incident count, policy violations, escalation rate	Governance quality
Learning	Reusable templates, improved workflows, and lessons captured	Organizational learning

5.6. Three-Phase AX Implementation Roadmap

The AX-5R framework can be translated into a three-phase roadmap: diagnose, redesign, and scale. This roadmap is designed for managers who need to move from an abstract AI strategy to a concrete implementation.

Phase 1 is diagnosis. The organization maps current AI usage, informal practices, candidate workflows, data sensitivity, high-stakes decisions, regulated activities, and accountability gaps. The output is an AX opportunity and risk map that separates low-risk, high-value opportunities from those that require stronger governance.

Phase 2 is redesign. The organization selects a small number of workflows. It specifies the current process, AI-supported process, review point, approval rule, data input rule, role level, and performance measure for each workflow. Training should cover workflow logic, data boundaries, review standards, escalation rules, and documentation, not only prompt writing.

Phase 3 is scaling, and it is a governed transfer process rather than replication. Scaling should occur only after baseline indicators show whether the redesigned workflow reduced cycle time, improved quality, reduced burden, increased stakeholder value, and avoided unacceptable risk. A validated pattern is then transferred through four adaptation steps. First, an eligibility test: the candidate workflow must belong to the same task family (for example, retrieval-summarize-draft-review), and its data sensitivity and decision stakes must be equal to or lower than those of the validated workflow; a higher-stakes candidate re-enters Phase 1 rather than inheriting validation. Second, risk-tier reclassification: controls are not copied but re-proportioned, since the same pattern may need source checks only in one workflow and named approval in another. Third, artifact adaptation: data input rules, role boundaries, escalation thresholds, and performance baselines are re-specified for the receiving workflow, with the receiving team, not the originating team, accountable for the adapted artifacts. Fourth, incident-lesson import: known failure modes and near misses from the validated workflow are written into the receiving workflow's controls before launch. A lightweight AX playbook preserves approved use cases, prohibited uses, workflow maps, role levels, risk controls, performance templates, and the lessons learned from accumulated incidents, making each subsequent transfer safer than the first.

The roadmap is summarized below.

In Table 12, the AX-5R framework is translated into a three-phase implementation roadmap. The table shows how managers can move from diagnosis to redesigned work and then to scaling, while keeping each phase tied to observable outputs rather than general AI enthusiasm.

Table 12. Phases, managerial tasks, key questions, and outputs of AX implementation.

Phase	Managerial Task	Key Questions	Output
Diagnose	Map AI use, workflows, and risks	Where is AI already used? Which workflows are suitable? What risks exist?	AX opportunity and risk map
Redesign	Rebuild selected workflows	What will AI do? Who reviews? What data are allowed? What measures matter?	Redesigned workflow and role protocol
Scale	Transfer validated patterns through eligibility testing, risk-tier reclassification, artifact adaptation, and incident-lesson import	Did the pilot improve outcomes? Is the receiving workflow in the same task family at equal or lower stakes? Which controls must be re-proportioned?	AX playbook, adaptation checklist, and scaling plan

This roadmap reinforces the central claim: AX begins with diagnosis of work, matures through redesigned workflows, and scales through measured outcomes, reusable artifacts, and governance routines.

5.7. Worked Workflow Vignettes

The following vignettes are illustrative application probes, not case-study evidence. They show how AX-5R converts AI enthusiasm into redesignable work systems.

In customer-service responses, adoption begins when agents use AI to draft replies. AX begins when the workflow is decomposed into intake, issue classification, policy retrieval, draft response, human review, and final send. AI supports retrieval and drafting; humans verify facts, tone, policy fit, and the need for escalation. Risk controls restrict

sensitive data, require source checks, and route high-risk complaints to supervisors. Return is measured through cycle time, first-contact resolution, correction rate, escalation rate, satisfaction, and reviewer burden. The transformation is not AI drafting itself, but the redesign of intake, retrieval, review, escalation, and measurement.

In internal management reporting, adoption begins when analysts use AI to summarize meeting notes, sales reports, or operating data. AX begins when the reporting workflow is redesigned around evidence traceability: data collection, retrieval, synthesis, draft generation, managerial review, and decision-log update. AI supports retrieval, summarization, and first drafts, while analysts remain accountable for source validation, interpretation, exceptions, and final recommendation. The return includes reporting time, factual corrections, decision-cycle speed, and template reuse. The transformation lies in evidence traceability and decision-log integration, not just faster report writing.

In HR document screening or compliance review, adoption begins when staff use AI to summarize resumes, policies, contracts, or training records. AX restricts AI to retrieval, summarization, and checklist preparation, while humans retain criteria interpretation, fairness review, final judgment, and exception handling. Risk controls require approved data handling, bias review, audit logs, approval gates, and prohibition of autonomous high-stakes decisions. Return includes review consistency, processing time, correction rates, workload, and compliance incidents. The transformation lies in restricting AI to assistive roles while preserving human judgment, fairness review, and auditability.

Across the three vignettes, the diagnostic questions remain stable: Is the workflow ready? What process changes? What role does AI play? Who remains accountable? What risks are controlled? What outcomes distinguish transformation from usage?

5.8. Propositions: Toward a Research Agenda for AX as Work System Transformation

The AX-5R framework is intended as a practical implementation model, but it can also support a focused research agenda. The propositions below translate the framework into six testable claims. They are deliberately stated at a level that can be examined through case studies, surveys, field experiments, longitudinal process research, or configurational analysis.

The six propositions are not equally urgent for initial validation. Empirical priority belongs to Propositions 2 and 6 for three reasons. First, they carry the framework's two load-bearing claims: redesign as the mediating mechanism and configuration as the unit of success. If either fails, the framework requires revision rather than refinement, making them the fastest routes to falsification. Second, both already have artifact-level initial evidence from the ablation probe (the largest single-component drop for redesign and the non-redundancy of the artifact-embedded components), so field studies can be designed to test a stated effect direction. Third, they discipline the remaining propositions: Proposition 1 specifies the entry condition that determines the admissible set of workflows in which the other propositions operate; Propositions 3 and 4 are moderation claims that presuppose the mediation structure of Proposition 2; and Proposition 5 specifies the feedback mechanism through which return evidence re-enters readiness and redesign. A practical sequence is therefore entry and eligibility (Proposition 1), mediation (Proposition 2, longitudinal or two-wave designs), configuration (Proposition 6, qualitative comparative analysis across workflows), moderators (Propositions 3 and 4, field experiments), and feedback (Proposition 5, throughout).

5.8.1. Workflow-Specific Readiness

Organizations often begin AI initiatives by selecting tools. However, the literature on digital transformation and IS assimilation suggests that access to tools does not guarantee organizational value [6,7,9–11]. Readiness is multidimensional and workflow-specific. It

includes leadership commitment, employee capability, data availability, workflow clarity, governance maturity, and measurement capability. Ecosystem-level analyses of AI value chains converge on the same point from the supply side: workforce readiness and managerial competency, rather than technical capacity, are the binding constraints on AI value realization [117]. A firm may be ready for internal document summarization but not ready for customer-facing advice, legal interpretation, or personnel evaluation.

Proposition 1. *Workflow-specific readiness, assessed at the workflow rather than solely at the firm or tool level, increases the likelihood that an adopted generative-AI use case enters the AX transformation cycle.*

Its downstream relationship with AX performance operates through the transformation core specified in Propositions 2 to 5.

5.8.2. Workflow Redesign as the Mediating Mechanism

Existing digital transformation and process-redesign research repeatedly shows that value arises when technologies reshape organizational processes [6,19,20,42,46]. Generative AI is no exception. AI tools may increase individual speed, but organizational transformation depends on whether the workflow itself changes. A simple binary adoption variable is therefore insufficient. Researchers should measure whether the process map changed, whether review points were added, whether tasks were reallocated, and whether performance indicators were revised.

Proposition 2. *Workflow redesign mediates the relationship between generative AI adoption and AX performance.*

5.8.3. Human–AI Role Clarity

AI transformation creates ambiguity because AI systems can generate outputs that appear to be work products. If AI drafts a report, recommends a vendor, summarizes a regulation, or classifies a customer request, responsibility can become unclear. Role clarity specifies what AI may do, what humans must verify, who approves final outputs, and which cases require escalation. This role logic is consistent with research on human–AI decision structures, reliance on automation, and conjoined agency [83,88,102,109].

Proposition 3. *Human–AI role clarity positively moderates the relationship between workflow redesign and AX performance by reducing accountability ambiguity, overreliance, and avoidance of useful AI support.*

5.8.4. Proportional Risk Governance

Governance is often treated as a constraint on innovation. In AX, however, appropriate oversight can enable scaling. Without oversight, managers may hesitate to expand AI use because risks remain unclear. With proportional oversight, firms can distinguish low-risk use cases that should be encouraged from high-risk use cases that should be restricted, consistent with risk-based and management-system approaches to AI governance [97,101,106,107].

Proposition 4. *Proportional risk governance strengthens the relationship between workflow redesign and AX performance by reducing uncontrolled exposure without blocking low-risk experimentation.*

5.8.5. Balanced Return Measurement

Many AI initiatives are justified by expected time savings, but time savings are only one part of AX value. A workflow may become faster but less accurate. A team may produce more documents but require more review. A customer support process may respond faster but escalate more complaints. Therefore, return must be measured across efficiency, quality, human, risk, customer, and learning indicators, consistent with balanced performance and IS-success measurement traditions [10,11,47].

Proposition 5. *Balanced AX performance measurement increases the likelihood that evidence from AI pilots feeds back into workflow readiness and redesign, enabling scalable organizational learning rather than isolated productivity anecdotes.*

5.8.6. Configurational Alignment and Strategic Consequence

The five components of AX-5R are mutually reinforcing, but they do not operate at the same level of analysis. Readiness primarily serves as a workflow entry and a precondition diagnosis, whereas redesign, role, risk, and return are more directly embedded in implementation artifacts. The five components fail in chained ways when developed in isolation, as elaborated in Section 6.1; this chaining is what makes AX performance configurational rather than additive: at the workflow or workflow-portfolio level, and conditional on meeting minimum readiness conditions, firms need alignment across the four transformation functions rather than excellence in one component alone.

Proposition 6. *Among workflows meeting minimum readiness conditions, AX performance is higher when redesign, role, risk, and return form an internally coherent configuration fitted to the workflow's task, risk, expertise, and regulatory conditions than when any one transformation function is developed in isolation.*

Alignment denotes coherent fit among the four transformation functions under the workflow's boundary conditions; it does not require equal component maturity or identical design weights across workflows.

Proposition 6 gains empirical content when the space of configurations is made explicit. Three archetypal configurations, defined by which components carry design weight under which boundary conditions, illustrate the testable structure. Table 13 states them as predictions, not findings.

Table 13. Three AX-5R configuration archetypes and a low-expertise corollary, with their boundary conditions (predictions under Proposition 6; the corollary row states the corresponding Proposition 3 failure prediction).

Archetype	Typical Boundary Conditions (Section 6.2)	Component Weighting	Predicted Trajectory	Predicted Failure Mode If Misapplied
Guarded pilot	High task risk or high data sensitivity; thin reviewer expertise; regulated sector	Readiness, Role, and Risk carry most weight; Redesign scope deliberately narrow; Return tracks risk and quality before efficiency	Slow but safe value; expertise accumulates before scale	Governance paralysis if Risk weighting is never relaxed after evidence accumulates

		All four transformation functions mutually reinforcing, with design weights fitted to the workflow's task, risk, expertise, and regulatory conditions		
Balanced transformer	Moderate task risk; adequate reviewer expertise; supportive culture	The reference AX path: pilot, measured validation, scaled playbook	Drift into archetype three if Return discipline weakens during scaling	
Velocity-first adopter	Low perceived risk; high tool readiness; weak governance tradition	Redesign and local Return emphasized; Role and Risk underdeveloped	Fast local gains that plateau; hidden exposure accumulates	Silent failure: usage grows while accountability, auditability, and learning do not
(Corollary)	Any archetype under sustained low expertise	Role approval authority formally present but substantively weak	Oversight becomes symbolic (Section 6.2, third condition)	Proposition 3's moderation fails; predicted null or negative redesign effect

The archetypes yield comparative predictions that configurational methods can test directly: velocity-first adopters should show the largest gap between usage metrics and transformation outcomes; guarded pilots should show the smallest incident rates but the longest time to value; and balanced transformers should dominate both on sustained AX performance. Configurations, not component scores, are the predicted unit of success, which is what distinguishes Proposition 6 from a maturity claim. The balanced-transformer row functions as the reference configuration implied by Proposition 6 itself; the discriminating predictions are carried by the two unbalanced archetypes and the corollary.

The strategic implication is deliberately framed as a research program claim rather than an empirical finding. The paper does not claim that AX-5R has already been field-validated as a source of durable competitive advantage; rather, it specifies a testable higher-order capability that future studies can examine to explain why firms with similar AI tools differ in AI value realization.

6. Discussion

6.1. Contributions to Systems-Based AI Transformation

AX-5R is not proposed as another maturity checklist. It specifies a minimum socio-technical systems architecture for AI-enabled work transformation: boundary and pre-condition diagnosis (readiness), task-interdependence restructuring (redesign), agency and accountability allocation (role), control and escalation loops (risk), and outcome feedback and learning (return). In doing so, AX-5R operationalizes the joint-optimization principle of socio-technical systems theory for the generative-AI context. Fragmented AI use becomes accountable, governable, and measurable transformation only when the technical artifact, work structure, human role system, governance controls, and feedback mechanisms are redesigned together.

The paper makes four cumulative increments. First, it defines the boundary between AI adoption and AX at the work-system level rather than treating transformation as tool use, firm-level aspiration, or governance compliance. Second, it specifies the mechanism of conversion: redesign mediates the path from adoption to AX performance, making Proposition 2 the process-theoretic core. Third, it specifies the systems condition of success: conditional on readiness, redesign, role, risk, and return must form an internally coherent configuration, making Proposition 6 the configurational core. Fourth, it renders these claims observable through workflow maps, role protocols, control artifacts, balanced dashboards, and a supplementary discriminant probe, so that the contribution is a

readiness-gated socio-technical architecture with an explicit mechanism, configuration logic, and observable implementation traces rather than another list of adoption factors.

The configurational claim in Proposition 6 is therefore, itself, a systems property. The four transformation functions are interdependent, and readiness gates their operation, so strengthening one in isolation leaves the larger work system underspecified; the chained failure modes named in Section 4.6 are the concrete form of this interdependence. AX performance therefore depends on the coherent configuration of these functions under a satisfied readiness gate rather than on excellence in any single component.

AX-5R operationalizes the joint-optimization principle of socio-technical systems theory in a specific way. The five Rs should not be assigned one-to-one to either the technical or the social subsystem, because doing so would reproduce the very separation that socio-technical theory warns against. Instead, each R is a coupling mechanism: it names an interface at which the technical and social subsystems must be jointly designed. Readiness tests whether technical feasibility and social absorptive and review capacity jointly permit a workflow to enter or re-enter the cycle. Redesign jointly restructures AI affordances and task interdependencies. Role couples system autonomy with human authority and accountability. Risk couples technical controls with organizational oversight and escalation. Return couples technical performance signals with social value and learning, feeding evidence back into reconfiguration and renewed readiness assessment. Table 14 states these five interfaces explicitly.

Table 14. AX-5R as a socio-technical coupling architecture: how each function jointly designs the technical and social subsystems.

AX-5R Function	Technical Subsystem	Social Subsystem	Coupling Function
Readiness (entry/re-entry gate)	Data quality and access, tool and model fit, infrastructure, security, traceability	Leadership commitment, employee capability, reviewer expertise, policy and measurement maturity	Tests whether technical feasibility and social capacity jointly permit a workflow to enter or re-enter the transformation cycle
Redesign	AI affordances, input, output, and evidence requirements, level of automation, system integration	Task sequence, routines, coordination, job design, review burden	Reconfigures task interdependencies jointly across AI capability and human work
Role	Authority scope, autonomy level, interface, handoff, override and logging	Decision rights, accountability, expertise, approval and exception handling	Allocates authority and handoffs between technical action and human judgment and responsibility
Risk	Data controls, validation, access control, logging, model monitoring, drift detection	Oversight rules, escalation authority, audit, legal and ethical accountability	Calibrates technical controls to the workflow's social stakes and risk ownership, proportional to stakes
Return	Accuracy, error, reliability, latency, drift, and operational signals	Quality, workload, customer value, capability, learning, strategic value	Integrates technical performance and social value signals into scale, revise, or stop decisions and renewed readiness assessment

Observable implementation artifacts are summarized in Table 5 and elaborated in Section 5.

In Table 14, AX-5R is recast in systems terms: the five components are not a checklist but the functional architecture, that is, diagnosis, restructuring, allocation, control, and feedback, that a work system requires in order to convert AI tool use into governed transformation.

This positioning locates AX-5R within systems and socio-technical traditions that hold that technology creates organizational value only when social and technical subsystems are jointly optimized and redesigned [20,24–26]. Generative AI sharpens this classic problem because adoption is now inexpensive, decentralized, cognitive, probabilistic, and often informal. The binding constraint on transformation has therefore shifted from tool access to the disciplined reconfiguration of the work system. AX-5R specifies the implementation architecture for meeting that constraint.

Read against the adjacent framework families in Table 3, the contributions can be stated comparatively. Where socio-technical and work system theory establish that joint optimization is necessary [24–26], AX-5R specifies the order and artifacts through which it is executed in generative AI. Where digital transformation research locates value in organizational change rather than deployment [5–8,113], AX-5R converts that thesis into workflow-level implementation moves. Where maturity models diagnose stages, AX-5R provides the mechanism for movement between stages. Where governance frameworks specify control obligations [97,101,106,107], AX-5R positions those controls inside the redesign cycle so that control and value creation are co-designed. The ablation probe adds artifact-level discriminant evidence that this integration is substantive rather than nominal: a sham five-part scaffold of equal form but no transformation content scored far lower than the full frame (Table 6).

6.2. Boundary Conditions

The AX-5R framework is designed for general business organizations, especially knowledge-intensive and administrative workflows, but its application depends on context. Several boundary conditions should be recognized.

First, task risk matters. Low-risk knowledge work, such as internal summarization, drafting, and document search, can often be performed earlier than high-stakes decision work. Legal, medical, financial, employment, safety, and public-sector decisions require stronger governance and may not be appropriate for early AX experimentation.

Second, data sensitivity matters. Workflows involving confidential, personal, regulated, or proprietary information require input controls and approved tools. A workflow may be technically suitable for AI but governance-ineligible because the data cannot be safely used.

Third, domain expertise matters, and the framework treats it as a designed input rather than an assumption. AI outputs can only be reviewed effectively when reviewers can judge output quality; where expertise is thin, human oversight becomes symbolic rather than substantive. AX-5R responds at three points. In readiness, reviewer capability is an explicit entry criterion: a workflow whose outputs people cannot evaluate is not ready for AI-enabled redesign at that autonomy level, regardless of tool quality. In role, the allocation matrix must name a competent approver, not merely an available one. Where none exists, the workflow is restricted to lower AI autonomy, sampled external review, or exclusion from automation. In return, correction and escalation rates function as a running indicator of whether the review is substantive. Expertise scarcity, therefore, narrows the admissible configuration space rather than silently weakening oversight, and it is the substantive condition behind the predicted weakening of Proposition 3's moderation in low-expertise contexts (Table 13).

Fourth, organizational culture matters. AX requires employees to experiment, report errors, share reusable practices, and accept redesigned roles. If employees fear replacement or punishment, AI use may become hidden, performative, or resisted.

Fifth, firm size matters. Large firms may have formal governance teams, legal departments, audit functions, and AI centers of excellence. Smaller firms may need simpler checklists and lightweight controls. AX-5R can be adapted to both, but the

implementation artifacts should differ in complexity. Integrative evidence on SME digital transformation supports this differentiated design expectation [113].

Sixth, the regulatory environment matters. Organizations in highly regulated sectors must align AX with legal and compliance requirements. In such contexts, role and risk should be developed before aggressive scaling, and some high-impact uses may need to remain outside automation.

These boundary conditions do not invalidate AX-5R; they parameterize the propositions. Task risk and data sensitivity shift the design weight toward role and risk, thereby setting the governance intensity at which Proposition 4's moderation should peak. Reviewer domain expertise is the substantive condition behind Proposition 3: where expertise is thin, role clarity on paper cannot deliver accountable review, so the moderation is predicted to weaken (Table 13's corollary row). Culture and firm size determine whether redesigned routines are reported and stabilized, which conditions the mediation path of Proposition 2. The regulatory environment fixes which configurations are even admissible, restricting the configuration space over which Proposition 6 operates. Each boundary condition is therefore not a caveat appended to the framework. Still, a moderator or scope condition is attached to a specific proposition, and Table 13 states the resulting configuration-level predictions.

6.3. Implications for Practice

The central practical implication is that AX is not a technology rollout. It is realized through an organizational execution system. Generative AI lowers the cost of adoption, but it does not lower the difficulty of transformation. In fact, easy adoption may increase the difficulty of transformation because AI use can spread before workflows, responsibilities, controls, and measures are ready.

AX-5R provides a way to manage this tension. Readiness prevents premature scaling. Redesign prevents superficial tool insertion. Role prevents accountability ambiguity. Risk prevents uncontrolled exposure. Return prevents unmeasured enthusiasm. Together, the five components translate scattered AI use into an organizational learning cycle.

The framework also avoids two symmetrical errors. The first is automation maximalism: applying AI wherever possible without adequate oversight. The second is governance paralysis: blocking useful low-risk applications because high-risk cases are difficult. AX requires differentiated governance. Low-risk use cases can be encouraged with light but explicit controls. High-risk use cases require strict human approval, auditability, escalation paths, or exclusion from automation.

For practitioners, the sequence is direct. Begin with a readiness diagnosis. Select a reviewable workflow. Map the current process. Insert AI only where it supports retrieval, summarization, drafting, classification, comparison, or recommendation. Define human review and approval. Add proportional risk controls. Measure baseline and post-pilot outcomes. Use performance, risk, and learning evidence to scale, revise, or stop the use case.

6.4. Future Research Agenda

The framework opens several research directions.

First, researchers can operationalize workflow-specific AX readiness. Current AI adoption studies often measure tool use or intention to use. Future studies should measure workflow clarity, governance maturity, data availability, review expertise, employee AI literacy, and measurement capability.

Second, researchers can study workflow redesign as a mediator. Empirical work should examine whether AI adoption improves performance directly or whether the effect

is mediated by process redesign. This would move AX research beyond simple adoption-performance models.

Third, researchers can examine human–AI role clarity. Field experiments could compare teams with identical AI tools but different role protocols. Outcome measures could include output quality, error rates, employee confidence, escalation behavior, and clarity of accountability.

Fourth, researchers can study proportional governance. Rather than asking whether governance helps or hinders innovation, studies can examine which governance intensity fits which risk level. This would produce a more nuanced AI governance theory.

Fifth, researchers can develop AX implementation maturity models. Such models should avoid generic maturity scoring and instead assess whether, once minimum readiness conditions are met, redesign, role, risk, and return are coherently aligned at the workflow or workflow-portfolio level.

Sixth, researchers can investigate industry-specific adaptations. The framework can be applied to healthcare, finance, education, manufacturing, logistics, professional services, and public administration. Each context will require different role boundaries, risk controls, auditability requirements, and return indicators.

Seventh, researchers can study the long-term strategic consequences of AX implementation capability. As generative AI tools become widely available, the advantage may shift from tool possession to organizational capability in redesigning and governing AI-enabled work.

Eighth, researchers can develop AI-augmented validation protocols for AX. Such protocols should not ask language models whether AX-5R is correct or treat LLM agreement as validation. Instead, they should use AI systems to generate controlled scenario variants or document-grounded workflow artifacts, then evaluate those artifacts through pre-specified human, expert, or rule-based adjudication. This would let researchers stress-test boundary conditions while avoiding the error of treating fluent AI output as empirical evidence.

The research agenda is summarized below.

In Table 15, the conceptual framework is converted into an empirical research agenda. The table identifies testable themes, example research questions, and possible methods to help AX-5R move from conceptual synthesis to cumulative empirical validation.

Table 15. Future Research Agenda for AX Implementation.

Research Theme	Example Question	Possible Method
Readiness	What dimensions predict successful AX scaling?	Survey, maturity assessment
Redesign	Does workflow redesign mediate AI adoption and performance?	Longitudinal field study
Role	How does role clarity affect trust and output quality?	Field experiment
Risk	What level of oversight best supports responsible scaling?	Comparative case study
Return	Which indicators predict sustained AX value?	Panel data, case comparison
Configuration	Which AX-5R configurations produce success?	QCA, process research
Strategy	Does AX implementation capability create an advantage?	Multi-firm comparative study
Validation architecture	Does AX-5R survive scenario, workflow, and expert-coding probes?	Scenario probe, document-grounded simulation, external coder validation

Taken together, these directions define AX not as a finished instrument to be accepted or rejected, but as a cumulative research program for diagnosing, redesigning, governing, measuring, and scaling AI-enabled work systems.

7. Conclusions

Generative AI adoption has become easier with decentralized access; however, AI transformation remains difficult because work systems must still be redesigned. This paper argued that the central challenge is not access to generative AI tools but the organizational ability to redesign work, allocate responsibility, govern risk, and measure outcomes. The proposed AX-5R framework integrates readiness, redesign, role, risk, and return into a socio-technical systems architecture for AI-enabled work system transformation.

The article contributes to systems-based AI transformation research by synthesizing seven literature streams into a single AX execution architecture and by explicitly deriving the failure modes of AX-5R. Its central novelty is to specify how fragmented AI tool use can be converted into redesigned, accountable, governable, and measurable organizational work systems. It contributes to practice by providing implementation artifacts, including diagnostic maps, role-allocation tools, risk controls, performance indicators, and workflow vignettes, that managers can use to move from pilots to scalable transformation.

The evidence boundary is explicit. The framework's derivation is conceptual and integrative; the ablation probe and expert-rating check provide artifact-level discriminant evidence for the non-redundancy of the artifact-embedded components; and organizational performance validation remains ahead, through the case, survey, experimental, configurational, and longitudinal designs specified in Section 5.8 and Table 15. Within that boundary, the framework already does the work this genre requires: it converts the fragmented literature into a single, falsifiable execution architecture with observable artifacts.

The practical message is direct. Organizations should not ask only whether employees are using generative AI. They should also ask whether AI use has changed the work system in a governed and measurable way. If the workflow is unchanged, the role boundary is unclear, the risk control is informal, and the return measure is anecdotal, the organization has adopted AI but has not undergone transformation. If those elements are redesigned together, adoption progresses to transformation. AX-5R specifies the minimum socio-technical architecture for executing, governing, and measuring that conversion.

Author Contributions: Conceptualization, K.S.S. and I.S.K.; methodology, K.S.S.; software, K.S.S.; validation, K.S.S. and I.S.K.; formal analysis, K.S.S.; investigation, K.S.S.; resources, K.S.S.; data curation, K.S.S.; writing of the original draft, K.S.S.; writing, review, and editing, K.S.S., I.S.K. and M.L.; visualization, K.S.S. and M.L.; supervision, K.S.S.; project administration, K.S.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: The ablation discriminant probe analyzed only language-model-generated artifacts and did not involve human participants or animals. In the expert-rating check (Appendix A.8), three domain experts served as professional evaluators of these artifacts rather than as research subjects: no intervention was performed; the raters' names and completion dates were recorded on the returned scoring sheets for administrative purposes only and were pseudonymized before analysis; no other personal or sensitive data were collected, no data about the raters themselves were analyzed, and their ratings are reported only in anonymized, aggregate form. Ethical review and approval were therefore not required in accordance with Article 2(1) of the Bioethics and Safety Act of the Republic of Korea and Article 2 of its Enforcement Rule, which define human subjects research as research involving physical intervention in a person, interaction to

collect data about a person, or the use of personally identifiable information for research purposes, none of which applies to this expert evaluation of study materials.

Informed Consent Statement: All three expert raters were informed in writing of the rating task, the blinding design, and the aggregate-only reporting of their ratings, and participated voluntarily on that basis.

Data Availability Statement: The data and analysis code for the ablation discriminant probe reported in Section 3 and Appendix A are openly available in Zenodo at <https://doi.org/10.5281/zenodo.20712342>. The data and documentation are released under a Creative Commons Attribution 4.0 license, and the analysis code under the MIT License. The archive contains the 252 generated artifacts, the independent cross-provider judge scores, the reported analysis output, the analysis scripts (which include the generation prompts and scoring rubric), and the expert-rating materials (rater instructions, the blinded artifact subsample, the reliability and convergence analysis script, and the aggregate expert-rating results). Individual expert ratings are not deposited, consistent with the aggregate-only reporting to which the raters agreed. The machine-scored results, including the effect sizes and the cross-provider score correlation ($r = 0.66$) in Table 6, are reproducible from the included data without API calls.

Acknowledgments: Use of Generative AI: Generative AI language models were used in two distinct capacities, disclosed here for transparency. First, as a study object and measurement instrument: the ablation discriminant probe (Section 3 and Appendix A) used gpt-4o-mini and claude-haiku-4-5-20251001 to generate the implementation artifacts that constitute the study's data, which were then scored by independent cross-provider model judges under a fixed rubric. Second, as a writing aid, generative AI tools assisted with language editing and formatting. All intellectual content, analytical decisions, interpretations, and conclusions are the authors' own. No AI system is listed as, or qualifies as, an author, consistent with the publisher's authorship policy. The authors reviewed, edited, and take full responsibility for all AI-assisted text and all interpretations reported in the manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Ablation Discriminant Probe Protocol

Appendix A.1. Purpose and Evidence Boundary

The ablation discriminant probe is a supplementary artifact-level test, an AI-augmented structural scenario probe. It does not use human subjects, observe firms, or measure organizational performance. It tests whether implementation artifacts generated under the full AX-5R frame receive higher implementation-quality ratings than artifacts generated under single-component ablation frames and a sham five-part frame of equal form but no transformation content. The unit of analysis is the generated artifact, not a person, team, firm, or field outcome. Accordingly, the probe is not treated as a field validation of AX-5R, but rather as a controlled artifact-level test to determine whether the framework produces distinguishable implementation outputs under matched task conditions.

Appendix A.2. Generators and Design

Two cross-provider language models served as generators: gpt-4o-mini and claude-haiku-4-5-20251001. Each generated artifacts for the three workflow contexts under seven arms, with six repetitions per generator-workflow-arm cell, producing 252 artifacts in total.

The same output constraints were used across conditions: 350 to 500 words, concise headings and bullets where helpful, enough operational detail for a manager to turn the output into a pilot checklist, no invented company names, no budgets, and no software vendor claims.

Appendix A.3. Workflow Contexts

The customer service workflow involved a mid-sized online retailer receiving product questions, refund requests, shipping complaints, and warranty claims. The management-reporting workflow concerned a business unit preparing monthly operating reports from spreadsheets, meeting notes, project trackers, and prior reports. The HR document-screening workflow covered promotion packets, job descriptions, training records, and employee development documents, with final employment-related decisions remaining human.

Appendix A.4. Arms

The full AX-5R arm asked for a plan built around readiness, redesign, role, risk, and return, working together. Each of the five ablation arms requested a plan built around four components, with one component removed and unnamed. The sham control used the PLANT method (Procure, Localize, Announce, Network, Tally). This plausible but transformation-irrelevant five-step structure does not encode work system redesign, role allocation, proportional oversight, or balanced return measurement. The sham controls whether any five-part structure, rather than AX-5R specifically, drives the result.

Appendix A.5. Independent Blinded Scoring

To reduce same-provider circularity and label-conformity bias, the scoring rubric evaluated implementation-quality criteria rather than explicit conformity to AX-5R labels, and each artifact was scored independently by both provider models. For the primary analysis, each artifact's score was taken from the judge belonging to the provider opposite to the generator: OpenAI-generated artifacts were scored by the Anthropic model, and Anthropic-generated artifacts were scored by the OpenAI model. Scores from both providers enabled estimation of the cross-provider score correlation.

Judges were blind to the arm and to any framework name. No second artifact was placed in the judge's context, eliminating pairwise position-order bias. Each judge scored five fixed implementation-quality criteria: workflow redesign depth, accountable human–AI role allocation, proportionate oversight, balanced measurable outcomes, and staged integration. Each criterion was scored from 0 to 2, and each artifact also received a holistic implementation-quality score from 0 to 10. The holistic score was the primary outcome. These criteria were selected because they represent generic implementation-quality dimensions derived from the literature synthesis, not because they require AX-5R labels to appear.

Scoring Criteria and Their Literature Basis

To address the concern that the rubric might simply reward conformity to AX-5R labels, Table A1 maps each scoring criterion to its independent literature basis. It shows that each rewards substantive implementation quality rather than the presence of an AX-5R term. The criteria were specified before scoring and do not name or require the AX-5R components.

Table A1. Scoring criteria, the literature basis, and distinction from AX-5R label conformity.

Scoring Criterion	Literature Basis	Distinction from Label Conformity
Workflow redesign depth	Work system theory; business process redesign; effective use	Rewards actual restructuring of task interdependencies, independent of AX-5R labels

Accountable human–AI role allocation	Human–AI collaboration; automation reliance; accountability	Rewards explicit approval and accountability boundaries, independent of AX-5R labels
Proportionate oversight	AI governance; risk management frameworks	Rewards risk-matched controls and escalation, independent of AX-5R labels
Balanced measurable outcomes	Information systems success; balanced scorecard; value realization	Rewards multidimensional outcome evidence, independent of AX-5R labels
Staged integration	Digital transformation; scaling; routinization	Rewards phased implementation sequencing, independent of AX-5R labels

Appendix A.6. Analysis

Because the two generators differed systematically in baseline verbosity and absolute score levels, holistic scores were standardized within the generator before pooling. This isolates the framing effect from generator-level score shifts. All contrasts were also estimated within each generator separately. For each arm, the full-versus-arm difference was summarized by Cohen's *d* with a 95% confidence interval and a Mann–Whitney test; one-sided tests were pre-specified for the directional discriminant expectation stated in advance, and all results for these contrasts are reported two-sided in Table 6 (Section 3.3), the more conservative convention.

These contrasts are interpreted as descriptive rather than confirmatory evidence of rating separation. The 36 artifacts per arm span three workflows and six repetitions per generator; the contrasts treat artifacts as independent after within-generator standardization and do not model workflow-level clustering, which is an additional reason the probe is reported as artifact-level evidence rather than confirmatory organizational evidence. As a context-level sensitivity check, treating the six generator-by-workflow cells as matched units and applying exact paired Wilcoxon tests to the per-cell arm means yields the same direction in all six cells for the sham, redesign, and risk contrasts and in five of six cells for the role and return contrasts; with only six matched units, however, no contrast can remain significant under Holm correction (all adjusted $p = 0.156$). The probe's inferential unit is therefore the artifact, the contrasts are described throughout as pre-specified artifact-level evidence rather than as confirmation of context-general effects, and extending the number of independent workflow and model contexts is part of the research agenda (Section 5.8). All five pre-specified contrasts (the sham control and the four artifact-relevant ablations) remained significant after Holm-Bonferroni correction for multiple comparisons under both the one-sided and the two-sided convention. The cross-provider score correlation was the Pearson correlation between the two providers' holistic scores across the 251 artifacts that received a valid score from both providers ($r = 0.66$); this indexes linear association across providers rather than absolute agreement. One artifact was excluded from this reliability estimate because one provider's response could not be parsed into a valid score, while the opposite provider's score used in the primary analysis remained valid.

Appendix A.7. Interpretation and Limits

The ablation discriminant probe should be interpreted within a narrow evidence boundary. It provides artifact-level evidence that the full AX-5R frame produces distinguishable implementation outputs under matched prompt conditions. It does not show that AX-5R improves organizational performance, predicts firm-level transformation outcomes, or has been validated in field settings.

The probe is also not equally sensitive to all five AX-5R components. Redesign, role, risk, and return are directly embedded in implementation artifacts, whereas readiness primarily serves as a workflow entry and an organizational precondition. Therefore, a null

readiness ablation should not be interpreted as evidence that readiness is unnecessary for AX. Rather, it indicates that an artifact-level probe is less sensitive to readiness than to artifact-embedded design features. Field studies, surveys, case research, and configurational analyses are needed to test whether AX-5R predicts transformation outcomes across organizations, workflows, industries, and regulatory settings.

Table A2. Per-arm descriptive statistics for the ablation discriminant probe (primary cross-provider judge scores).

Arm	<i>n</i>	Raw M (SD), Anthropic-Generated	Raw M (SD), OpenAI-Generated	Pooled Standardized M (SD)	95% CI of Standardized Mean
Full AX-5R	36	8.33 (0.69), <i>n</i> = 18	5.89 (0.47), <i>n</i> = 18	0.525 (0.790)	[0.258, 0.792]
Readiness removed	36	8.50 (0.71), <i>n</i> = 18	6.11 (0.47), <i>n</i> = 18	0.746 (0.807)	[0.473, 1.020]
Redesign removed	36	8.00 (0.00), <i>n</i> = 18	5.06 (0.54), <i>n</i> = 18	−0.082 (0.318)	[−0.190, 0.026]
Role removed	36	8.22 (0.55), <i>n</i> = 18	5.33 (0.69), <i>n</i> = 18	0.205 (0.711)	[−0.035, 0.446]
Risk removed	36	8.00 (0.00), <i>n</i> = 18	5.44 (0.62), <i>n</i> = 18	0.082 (0.406)	[−0.055, 0.220]
Return removed	36	8.17 (0.38), <i>n</i> = 18	5.39 (0.70), <i>n</i> = 18	0.186 (0.580)	[−0.010, 0.383]
Sham five-part control	36	7.22 (0.81), <i>n</i> = 18	2.72 (0.46), <i>n</i> = 18	−1.663 (0.981)	[−1.995, −1.332]

Raw scores are shown separately by generator because the two generators differ in baseline score levels, which is why scores were standardized within each generator before pooling (Section 3.3); the confidence intervals are for the pooled standardized means.

Appendix A.8. Human-Expert-Rating Check

To address the possibility that model judges share systematic biases that human experts would not, three doctoral-level experts, spanning senior management practice, public policy research, and university faculty, none of whom is an author, rated a stratified subsample of 56 artifacts (8 per arm; 28 per generator; the three workflows near-balanced within each arm) using the same five anchored criteria and holistic 0 to 10 scale as the model judges, blind to arm labels, framework identity, and study hypotheses. Raters worked independently in randomized presentation orders, and the materials were provided in each rater's language with identical rubric content. The convergence result is direct: the mean expert rating correlated $r = 0.78$ (Spearman $\rho = 0.78$, $n = 56$) with the cross-provider machine scores on the same artifacts, and the expert panel reproduced the discriminant core of the probe, rating the full frame far above the sham five-part control ($d = 1.90$, exact two-sided Mann–Whitney $p = 0.005$) and, mirroring the machine result, rating the Readiness-removed arm nominally highest. Inter-rater reliability for the holistic score was $ICC(2,1) = 0.29$ for a single rater and $ICC(2,3) = 0.55$ for the averaged rating used in the convergence analysis (consistency forms 0.34 and 0.61); quadratic-weighted kappa for the five ordinal criteria, averaged across rater pairs, ranged from −0.03 (integration) to 0.43 (role). The four artifact-ablation contrasts were not individually resolvable in this eight-per-arm subsample, which has power only for large effects; the primary artifact-level evidence for those contrasts remains the full-sample cross-provider panel ($n = 36$ per arm), and the expert check functions as a judge-validity audit rather than a replication of the experiment. The rating materials, the blinded artifact subsample, the analysis script, and the aggregate rating results are included in the public archive.

References

1. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the opportunities and risks of foundation models. *arXiv* **2021**, arXiv:2108.07258. <https://doi.org/10.48550/arXiv.2108.07258>.
2. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 1877–1901. <https://doi.org/10.48550/arXiv.2005.14165>.
3. Dwivedi, Y.K.; Kshetri, N.; Hughes, L.; Slade, E.L.; Jeyaraj, A.; Kar, A.K.; Baabdullah, A.M.; Koohang, A.; Raghavan, V.; Ahuja, M.; et al. So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *Int. J. Inf. Manag.* **2023**, *71*, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>.
4. Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F.L. GPT-4 technical report. *arXiv* **2023**, arXiv:2303.08774. <https://doi.org/10.48550/arXiv.2303.08774>.
5. Bharadwaj, A.; El Sawy, O.A.; Pavlou, P.A.; Venkatraman, N. Digital business strategy: Toward a next generation of insights. *MIS Q.* **2013**, *37*, 471–482. <https://doi.org/10.25300/misq/2013/37:2.3>.
6. Vial, G. Understanding digital transformation: A review and a research agenda. *J. Strateg. Inf. Syst.* **2019**, *28*, 118–144. <https://doi.org/10.1016/j.jsis.2019.01.003>.
7. Wessel, L.; Baiyere, A.; Ologeanu-Taddei, R.; Cha, J.; Blegind-Jensen, T. Unpacking the difference between digital transformation and IT-enabled organizational transformation. *J. Assoc. Inf. Syst.* **2021**, *22*, 102–129. <https://doi.org/10.17705/1jais.00655>.
8. Westerman, G.; Bonnet, D.; McAfee, A. *Leading Digital: Turning Technology into Business Transformation*; Harvard Business Review Press: Brighton, MA, USA, 2014.
9. Burton-Jones, A.; Grange, C. From use to effective use: A representation theory perspective. *Inf. Syst. Res.* **2013**, *24*, 632–658. <https://doi.org/10.1287/isre.1120.0444>.
10. DeLone, W.H.; McLean, E.R. Information systems success: The quest for the dependent variable. *Inf. Syst. Res.* **1992**, *3*, 60–95. <https://doi.org/10.1287/isre.3.1.60>.
11. DeLone, W.H.; McLean, E.R. The DeLone and McLean model of information systems success: A ten-year update. *J. Manag. Inf. Syst.* **2003**, *19*, 9–30. <https://doi.org/10.1080/07421222.2003.11045748>.
12. Dell’Acqua, F.; McFowland, E., III; Mollick, E.R.; Lifshitz-Assaf, H.; Kellogg, K.C.; Rajendran, S.; Krayner, L.; Candelon, F.; Lakhani, K.R. Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organ. Sci.* **2026**, *37*, 403–423. <https://doi.org/10.1287/orsc.2025.21838>.
13. Noy, S.; Zhang, W. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* **2023**, *381*, 187–192. <https://doi.org/10.1126/science.adh2586>.
14. Ajzen, I. The theory of planned behavior. *Organ. Behav. Hum. Decis. Process.* **1991**, *50*, 179–211. [https://doi.org/10.1016/0749-5978\(91\)90020-t](https://doi.org/10.1016/0749-5978(91)90020-t).
15. Barley, S.R. Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Adm. Sci. Q.* **1986**, *31*, 78–108. <https://doi.org/10.2307/2392767>.
16. Davis, F.D. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Q.* **1989**, *13*, 319–340. <https://doi.org/10.2307/249008>.
17. Markus, M.L. Power, politics, and MIS implementation. *Commun. ACM* **1983**, *26*, 430–444. <https://doi.org/10.1145/358141.358148>.
18. Markus, M.L.; Robey, D. Information technology and organizational change: Causal structure in theory and research. *Manag. Sci.* **1988**, *34*, 583–598. <https://doi.org/10.1287/mnsc.34.5.583>.
19. Orlikowski, W.J. The duality of technology: Rethinking the concept of technology in organizations. *Organ. Sci.* **1992**, *3*, 398–427. <https://doi.org/10.1287/orsc.3.3.398>.
20. Orlikowski, W.J. Using technology and constituting structures: A practice lens for studying technology in organizations. *Organ. Sci.* **2000**, *11*, 404–428. <https://doi.org/10.1287/orsc.11.4.404.14600>.
21. Orlikowski, W.J.; Scott, S.V. Sociomateriality: Challenging the separation of technology, work and organization. *Acad. Manag. Ann.* **2008**, *2*, 433–474. <https://doi.org/10.5465/19416520802211644>.
22. Venkatesh, V.; Morris, M.G.; Davis, G.B.; Davis, F.D. User acceptance of information technology: Toward a unified view. *MIS Q.* **2003**, *27*, 425–478. <https://doi.org/10.2307/30036540>.
23. Zuboff, S. *The Age of the Smart Machine: The Future of Work and Power*; Basic Books: New York, NY, USA, 1988.

24. Trist, E.L.; Bamforth, K.W. Some social and psychological consequences of the longwall method of coal-getting. *Hum. Relat.* **1951**, *4*, 3–38. <https://doi.org/10.1177/001872675100400101>.
25. Bostrom, R.P.; Heinen, J.S. MIS problems and failures: A socio-technical perspective. *MIS Q.* **1977**, *1*, 17–32. <https://doi.org/10.2307/248710>.
26. Alter, S. Work system theory: Overview of core concepts, extensions, and challenges for the future. *J. Assoc. Inf. Syst.* **2013**, *14*, 72–121. <https://doi.org/10.17705/1jais.00323>.
27. Alves, M.; Martinho, D.; Marcão, R.; Sobreiro, P. Generative AI adoption in B2B firms: Ethical governance, innovation capabilities, and long-term competitive performance. *Systems* **2026**, *14*, 410. <https://doi.org/10.3390/systems14040410>.
28. Zhang, Y.; Dong, C. Exploring the digital transformation of generative AI-assisted foreign language education: A socio-technical systems perspective based on mixed-methods. *Systems* **2024**, *12*, 462. <https://doi.org/10.3390/systems12110462>.
29. Mezher, M.A.; Gunawan, I.; Fayezi, S. Lean 4.0 as a socio-technical system: Mapping the interaction of soft practices and Industry 4.0 in digital transformation. *Systems* **2026**, *14*, 9. <https://doi.org/10.3390/systems14010009>.
30. Kane, G.C.; Palmer, D.; Phillips, A.N.; Kiron, D.; Buckley, N. *Strategy, Not Technology, Drives Digital Transformation*; MIT Sloan Management Review and Deloitte University Press: Cambridge, MA, USA, 2015. Available online: <https://sloanreview.mit.edu/projects/strategy-drives-digital-transformation/> (accessed on 10 June 2026).
31. Kotter, J.P. Leading change: Why transformation efforts fail. *Harv. Bus. Rev.* **1995**, *73*, 59–67. Available online: <https://hbr.org/1995/05/leading-change-why-transformation-efforts-fail-2> (accessed on 10 June 2026).
32. Li, F. Leading digital transformation: Three emerging approaches for managing the transition. *Int. J. Oper. Prod. Manag.* **2020**, *40*, 809–817. <https://doi.org/10.1108/ijopm-04-2020-0202>.
33. Lucas, H.C., Jr.; Goh, J.M. Disruptive technology: How Kodak missed the digital photography revolution. *J. Strateg. Inf. Syst.* **2009**, *18*, 46–55. <https://doi.org/10.1016/j.jsis.2009.01.002>.
34. Matt, C.; Hess, T.; Benlian, A. Digital transformation strategies. *Bus. Inf. Syst. Eng.* **2015**, *57*, 339–343. <https://doi.org/10.1007/s12599-015-0401-5>.
35. Matt, C.; Hess, T.; Benlian, A.; Wiesboeck, F. Options for formulating a digital transformation strategy. *MIS Q. Exec.* **2016**, *15*, 6. Available online: <https://aisel.aisnet.org/misqe/vol15/iss2/6/> (accessed on 10 June 2026).
36. Nambisan, S.; Lyytinen, K.; Majchrzak, A.; Song, M. Digital innovation management: Reinventing innovation management research in a digital world. *MIS Q.* **2017**, *41*, 223–238. <https://doi.org/10.25300/misq/2017/41:1.03>.
37. Romer, P.M. Endogenous technological change. *J. Political Econ.* **1990**, *98*, S71–S102. <https://doi.org/10.1086/261725>.
38. Sebastian, I.M.; Ross, J.W.; Beath, C.; Mockler, M.; Moloney, K.G.; Fonstad, N.O. How big old companies navigate digital transformation. *MIS Q. Exec.* **2017**, *16*, 6. Available online: <https://aisel.aisnet.org/misqe/vol16/iss3/6/> (accessed on 10 June 2026).
39. Warner, K.S.R.; Wager, M. Building dynamic capabilities for digital transformation: An ongoing process of strategic renewal. *Long. Range Plan.* **2019**, *52*, 326–349. <https://doi.org/10.1016/j.lrp.2018.12.001>.
40. Yoo, Y.; Henfridsson, O.; Lyytinen, K. The new organizing logic of digital innovation: An agenda for information systems research. *Inf. Syst. Res.* **2010**, *21*, 724–735. <https://doi.org/10.1287/isre.1100.0322>.
41. Cohen, W.M.; Levinthal, D.A. Absorptive capacity: A new perspective on learning and innovation. *Adm. Sci. Q.* **1990**, *35*, 128–152. <https://doi.org/10.2307/2393553>.
42. Davenport, T.H. *Process Innovation: Reengineering Work Through Information Technology*; Harvard Business School Press: Brighton, MA, USA, 1993.
43. Davenport, T.H.; Harris, J.G. *Competing on Analytics: The New Science of Winning*; Harvard Business School Press: Brighton, MA, USA, 2007.
44. Felin, T.; Foss, N.J.; Heimeriks, K.H.; Madsen, T.L. Microfoundations of routines and capabilities: Individuals, processes, and structure. *J. Manag. Stud.* **2012**, *49*, 1351–1374. <https://doi.org/10.1111/j.1467-6486.2012.01052.x>.
45. Grant, R.M. Toward a knowledge-based theory of the firm. *Strateg. Manag. J.* **1996**, *17*, 109–122. <https://doi.org/10.1002/smj.4250171110>.
46. Hammer, M.; Champy, J. *Reengineering the Corporation: A Manifesto for Business Revolution*; Harper Business: New York, NY, USA, 1993.
47. Kaplan, R.S.; Norton, D.P. The balanced scorecard: Measures that drive performance. *Harv. Bus. Rev.* **1992**, *70*, 71–79. Available online: <https://hbr.org/1992/01/the-balanced-scorecard-measures-that-drive-performance-2> (accessed on 10 June 2026).
48. Leonardi, P.M. When flexible routines meet flexible technologies: Affordance, constraint, and the imbrication of human and material agencies. *MIS Q.* **2011**, *35*, 147–167. <https://doi.org/10.2307/23043493>.
49. March, J.G. Exploration and exploitation in organizational learning. *Organ. Sci.* **1991**, *2*, 71–87. <https://doi.org/10.1287/orsc.2.1.71>.

50. Nonaka, I. A dynamic theory of organizational knowledge creation. *Organ. Sci.* **1994**, *5*, 14–37. <https://doi.org/10.1287/orsc.5.1.14>.
51. Nonaka, I.; Takeuchi, H. *The Knowledge-Creating Company*; Oxford University Press: Oxford, UK, 1995.
52. Teece, D.J. Explicating dynamic capabilities: The nature and microfoundations of sustainable enterprise performance. *Strateg. Manag. J.* **2007**, *28*, 1319–1350. <https://doi.org/10.1002/smj.640>.
53. Teece, D.J.; Pisano, G.; Shuen, A. Dynamic capabilities and strategic management. *Strateg. Manag. J.* **1997**, *18*, 509–533. [https://doi.org/10.1002/\(sici\)1097-0266\(199708\)18:7<509::aid-smj882>3.0.co;2-z](https://doi.org/10.1002/(sici)1097-0266(199708)18:7<509::aid-smj882>3.0.co;2-z).
54. Tushman, M.L.; O'Reilly, C.A. Ambidextrous organizations: Managing evolutionary and revolutionary change. *Calif. Manag. Rev.* **1996**, *38*, 8–30. <https://doi.org/10.2307/41165852>.
55. Zahra, S.A.; George, G. Absorptive capacity: A review, reconceptualization, and extension. *Acad. Manag. Rev.* **2002**, *27*, 185–203. <https://doi.org/10.2307/4134351>.
56. Acemoglu, D. The simple macroeconomics of AI. *Econ. Policy* **2025**, *40*, 13–58. <https://doi.org/10.1093/epolic/eiae042>.
57. Acemoglu, D.; Restrepo, P. *Artificial Intelligence, Automation and Work*; NBER Working Paper No. 24196; National Bureau of Economic Research: Cambridge, MA, USA, 2018. <https://doi.org/10.3386/w24196>.
58. Acemoglu, D.; Restrepo, P. Automation and new tasks: How technology displaces and reinstates labor. *J. Econ. Perspect.* **2019**, *33*, 3–30. <https://doi.org/10.1257/jep.33.2.3>.
59. Acemoglu, D.; Restrepo, P. Tasks, automation, and the rise in U.S. wage inequality. *Econometrica* **2022**, *90*, 1973–2016. <https://doi.org/10.3982/ecta19815>.
60. Agrawal, A.; Gans, J.; Goldfarb, A. Artificial intelligence: The ambiguous labor market impact of automating prediction. *J. Econ. Perspect.* **2019**, *33*, 31–50. <https://doi.org/10.1257/jep.33.2.31>.
61. Autor, D. *Applying AI to Rebuild Middle Class Jobs*; NBER Working Paper No. 32140; National Bureau of Economic Research: Cambridge, MA, USA, 2024. <https://doi.org/10.3386/w32140>.
62. Brynjolfsson, E.; Hitt, L.M.; Yang, S. Intangible assets: Computers and organizational capital. *Brook. Pap. Econ. Act.* **2002**, *2002*, 137–181. <https://doi.org/10.1353/eca.2002.0003>.
63. Brynjolfsson, E.; Li, D.; Raymond, L.R. *Generative AI at Work*; NBER Working Paper No. 31161; National Bureau of Economic Research: Cambridge, MA, USA, 2023. <https://doi.org/10.3386/w31161>.
64. Brynjolfsson, E.; McAfee, A. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*; W. W. Norton: New York, NY, USA, 2014.
65. Furman, J.; Seamans, R. AI and the economy. *Innov. Policy Econ.* **2019**, *19*, 161–191. <https://doi.org/10.1086/699936>.
66. Anthony, C.; Bechky, B.A.; Fayard, A.-L. Collaborating" with AI: Taking a system view to explore the future of work. *Organ. Sci.* **2023**, *34*, 1672–1694. <https://doi.org/10.1287/orsc.2022.1651>.
67. Bailey, D.E.; Faraj, S.; Hinds, P.J.; Leonardi, P.M.; von Krogh, G. We are all theorists of technology now: A relational perspective on emerging technology and organizing. *Organ. Sci.* **2022**, *33*, 1–18. <https://doi.org/10.1287/orsc.2021.1562>.
68. Baird, A.; Maruping, L.M. The next generation of research on IS use: A theoretical framework of delegation to and from agentic IS artifacts. *MIS Q.* **2021**, *45*, 315–341. <https://doi.org/10.25300/MISQ/2021/15882>.
69. Banks, S.; Formosa, P. The ethical implications of artificial intelligence (AI) for meaningful work. *J. Bus. Ethics* **2023**, *185*, 725–740. <https://doi.org/10.1007/s10551-023-05339-7>.
70. Berente, N.; Gu, B.; Recker, J.; Santhanam, R. Managing artificial intelligence. *MIS Q.* **2021**, *45*, 1433–1450. <https://doi.org/10.25300/MISQ/2021/16274>.
71. Choudhury, P.; Froughi, C.; Larson, B. Work-from-anywhere: The productivity effects of geographic flexibility. *Strateg. Manag. J.* **2021**, *42*, 655–683. <https://doi.org/10.1002/smj.3251>.
72. Davenport, T.H.; Kirby, J. *Only Humans Need Apply: Winners and Losers in the Age of Smart Machines*; Harper Business: New York, NY, USA, 2016.
73. Davenport, T.H.; Ronanki, R. Artificial intelligence for the real world. *Harv. Bus. Rev.* **2018**, *96*, 108–116. Available online: <https://hbr.org/2018/01/artificial-intelligence-for-the-real-world> (accessed on 10 June 2026).
74. Dellermann, D.; Ebel, P.; Söllner, M.; Leimeister, J.M. Hybrid intelligence. *Bus. Inf. Syst. Eng.* **2019**, *61*, 637–643. <https://doi.org/10.1007/s12599-019-00595-2>.
75. Faraj, S.; Pachidi, S.; Sayegh, K. Working and organizing in the age of the learning algorithm. *Inf. Organ.* **2018**, *28*, 62–70. <https://doi.org/10.1016/j.infoandorg.2018.02.005>.
76. Gregory, R.W.; Henfridsson, O.; Kaganer, E.; Kyriakou, H. The role of artificial intelligence and data network effects for creating user value. *Acad. Manag. Rev.* **2021**, *46*, 534–551. <https://doi.org/10.5465/amr.2019.0178>.

77. Huang, M.-H.; Rust, R.T. Artificial intelligence in service. *J. Serv. Res.* **2018**, *21*, 155–172. <https://doi.org/10.1177/1094670517752459>.
78. Iansiti, M.; Lakhani, K.R. *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*; Harvard Business Review Press: Brighton, MA, USA, 2020.
79. Jarrahi, M.H. Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Bus. Horiz.* **2018**, *61*, 577–586. <https://doi.org/10.1016/j.bushor.2018.03.007>.
80. Kellogg, K.C.; Valentine, M.A.; Christin, A. Algorithms at work: The new contested terrain of control. *Acad. Manag. Ann.* **2020**, *14*, 366–410. <https://doi.org/10.5465/annals.2018.0174>.
81. Kolbjørnsrud, V. Designing the intelligent organization: Six principles for human-AI collaboration. *Calif. Manag. Rev.* **2024**, *66*, 44–64. <https://doi.org/10.1177/00081256231211020>.
82. Krakowski, S.; Luger, J.; Raisch, S. Artificial intelligence and the changing sources of competitive advantage. *Strateg. Manag. J.* **2023**, *44*, 1425–1452. <https://doi.org/10.1002/smj.3387>.
83. Murray, A.; Rhymer, J.; Sirmon, D.G. Humans and technology: Forms of conjoined agency in organizations. *Acad. Manag. Rev.* **2021**, *46*, 552–571. <https://doi.org/10.5465/amr.2019.0186>.
84. Puranam, P. Human-AI collaborative decision-making as an organization design problem. *J. Organ. Des.* **2021**, *10*, 75–80. <https://doi.org/10.1007/s41469-021-00095-2>.
85. Raisch, S.; Fomina, K. Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Acad. Manag. Rev.* **2025**, *50*, 441–464. <https://doi.org/10.5465/amr.2021.0421>.
86. Raisch, S.; Krakowski, S. Artificial intelligence and management: The automation-augmentation paradox. *Acad. Manag. Rev.* **2021**, *46*, 192–210. <https://doi.org/10.5465/amr.2018.0072>.
87. Seeber, I.; Bittner, E.A.C.; Briggs, R.O.; de Vreede, T.; de Vreede, G.-J.; Elkins, A.; Maier, R.; Merz, A.B.; Oeste-Reiß, S.; Randrup, N.L.; et al. Machines as teammates: A research agenda on AI in team collaboration. *Inf. Manag.* **2020**, *57*, 103174. <https://doi.org/10.1016/j.im.2019.103174>.
88. Shrestha, Y.R.; Ben-Menahem, S.M.; von Krogh, G. Organizational decision-making structures in the age of artificial intelligence. *Calif. Manag. Rev.* **2019**, *61*, 66–83. <https://doi.org/10.1177/0008125619862257>.
89. Tambe, P.; Cappelli, P.; Yakubovich, V. Artificial intelligence in human resources management: Challenges and a path forward. *Calif. Manag. Rev.* **2019**, *61*, 15–42. <https://doi.org/10.1177/0008125619867910>.
90. van der Aalst, W.M.P.; Bichler, M.; Heinzl, A. Robotic process automation. *Bus. Inf. Syst. Eng.* **2018**, *60*, 269–272. <https://doi.org/10.1007/s12599-018-0542-4>.
91. Verganti, R.; Vendraminelli, L.; Iansiti, M. Innovation and design in the age of artificial intelligence. *J. Product. Innov. Manag.* **2020**, *37*, 212–227. <https://doi.org/10.1111/jpim.12523>.
92. von Krogh, G. Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Acad. Manag. Discov.* **2018**, *4*, 404–409. <https://doi.org/10.5465/amd.2018.0084>.
93. Amershi, S.; Weld, D.; Vorvoreanu, M.; Fournier, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P.N.; Inkpen, K.; et al. Guidelines for human-AI interaction. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019. <https://doi.org/10.1145/3290605.3300233>.
94. Arrieta, A.B.; Diaz-Rodriguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; Garcia, S.; Gil-Lopez, S.; Molina, D.; Benjamins, R.; et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **2020**, *58*, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>.
95. Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y.T.; Li, Y.; Lundberg, S.; et al. Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv* **2023**, arXiv:2303.12712. <https://doi.org/10.48550/arXiv.2303.12712>.
96. Dietvorst, B.J.; Simmons, J.P.; Massey, C. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *J. Exp. Psychol. General.* **2015**, *144*, 114–126. <https://doi.org/10.1037/xge0000033>.
97. European Parliament and Council of the European Union. *Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*; European Parliament and Council of the European Union: Brussels, Belgium, 2024. Available online: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj> (accessed on 10 June 2026).
98. Floridi, L.; Chiriatti, M. GPT-3: Its nature, scope, limits, and consequences. *Minds Mach.* **2020**, *30*, 681–694. <https://doi.org/10.1007/s11023-020-09548-1>.
99. Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J.W.; Wallach, H.; Daume, H., III; Crawford, K. Datasheets for datasets. *Commun. ACM* **2021**, *64*, 86–92. <https://doi.org/10.1145/3458723>.

100. Glikson, E.; Woolley, A.W. Human trust in artificial intelligence: Review of empirical research. *Acad. Manag. Ann.* **2020**, *14*, 627–660. <https://doi.org/10.5465/annals.2018.0057>.
101. ISO/IEC 42001:2023; Information Technology—Artificial Intelligence—Management System. International Organization for Standardization: Geneva, Switzerland, 2023. Available online: <https://www.iso.org/standard/42001> (accessed on 10 June 2026).
102. Lee, J.D.; See, K.A. Trust in automation: Designing for appropriate reliance. *Hum. Factors* **2004**, *46*, 50–80. <https://doi.org/10.1518/hfes.46.1.50.30392>.
103. Liang, P.; Bommasani, R.; Lee, T.; Tsipras, D.; Soylu, D.; Yasunaga, M.; Zhang, Y.; Narayanan, D.; Wu, Y.; Kumar, A.; et al. Holistic evaluation of language models. *arXiv* **2022**, arXiv:2211.09110. <https://doi.org/10.48550/arXiv.2211.09110>.
104. McKinsey & Company. *The State of AI in 2025: Agents, Innovation, and Transformation*; McKinsey & Company: New York, NY, USA, 2025. Available online: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai> (accessed on 10 June 2026).
105. Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I.D.; Gebru, T. Model cards for model reporting. In Proceedings of the Conference on Fairness, Accountability, and Transparency, Atlanta, GA, USA, 29–31 January 2019; pp. 220–229. <https://doi.org/10.1145/3287560.3287596>.
106. NIST AI 100-1; Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology: Gaithersburg, MD, USA, 2023. <https://doi.org/10.6028/NIST.AI.100-1>.
107. NIST AI 600-1; Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. <https://doi.org/10.6028/NIST.AI.600-1>.
108. OECD. *Recommendation of the Council on Artificial Intelligence*; OECD: Berlin, Germany, 2019. Available online: <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449> (accessed on 10 June 2026).
109. Parasuraman, R.; Sheridan, T.B.; Wickens, C.D. A model for types and levels of human interaction with automation. *IEEE Trans. Syst. Man. Cybern.-Part A Syst. Hum.* **2000**, *30*, 286–297. <https://doi.org/10.1109/3468.844354>.
110. Raji, I.D.; Smart, A.; White, R.N.; Mitchell, M.; Gebru, T.; Hutchinson, B.; Smith-Loud, J.; Theron, D.; Barnes, P. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, Barcelona, Spain, 27–30 January 2020; pp. 33–44. <https://doi.org/10.1145/3351095.3372873>.
111. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 5998–6008. Available online: <https://papers.neurips.cc/paper/7181-attention-is-all-you-need> (accessed on 10 June 2026).
112. Weidinger, L.; Mellor, J.; Rauh, M.; Griffin, C.; Uesato, J.; Huang, P.S.; Cheng, M.; Glaese, A.; Balle, B.; Kasirzadeh, A.; et al. Ethical and social risks of harm from language models. *arXiv* **2021**, arXiv:2112.04359. <https://doi.org/10.48550/arXiv.2112.04359>.
113. Ben Slimane, S.; Coeurderoy, R.; Mhenni, H. Digital transformation of small and medium enterprises: A systematic literature review and an integrative framework. *Int. Stud. Manag. Organ.* **2022**, *52*, 96–120. <https://doi.org/10.1080/00208825.2022.2072067>.
114. Cheng, M.; Chong, H.-Y.; Xu, Y.; Wu, H. Exploring the determinants of generative artificial intelligence use intention: A mixed methods study. *J. Constr. Eng. Manag.* **2026**, *152*, 04026091. <https://doi.org/10.1061/JCEMD4.COENG-17763>.
115. Mogahed, M.; Mansouri, M. Towards governance of socio-technical system of systems: Leveraging lessons from proven engineering principles. *Systems* **2025**, *13*, 1113. <https://doi.org/10.3390/systems13121113>.
116. Sadiq, R.B.; Safie, N.; Abd Rahman, A.H.; Goudarzi, S. Artificial intelligence maturity model: A systematic literature review. *PeerJ Comput. Sci.* **2021**, *7*, e661. <https://doi.org/10.7717/peerj-cs.661>.
117. Billones, R.K.C.; Lauresta, D.A.S.; Dellosa, J.T.; Bong, Y.; Stergioulas, L.K.; Yunus, S. AI ecosystem and value chain: A multi-layered framework for analyzing supply, value creation, and delivery mechanisms. *Technologies* **2025**, *13*, 421. <https://doi.org/10.3390/technologies13090421>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.