

Artificial Intelligence and the Future of Work: Towards a Framework for Human-AI Augmentation, Workforce Resilience, and Equitable Transition

Piu Ghosh

Apple Supply Chain & CapEx Management, Cupertino, California
piu_ghosh@apple.com

Received: Jul 28, 2026, **Revised:** Aug 25, 2026, **Accepted:** Aug 28, 2026, **Published Online:** Aug 29, 2026

Reviewers: Anonymous Peer Review

Citation: Ghosh, P. (2026). Artificial Intelligence and the Future of Work: Towards a Framework for Human-AI Augmentation, Workforce Resilience, and Equitable Transition. *International Journal of Supply Chain Management*, 15(4), 7-23, <https://doi.org/10.59160/ijscm.v15i4.6418>

Abstract — This paper introduces two novel analytical frameworks—the Human-AI Collaboration Spectrum and the Work Transition Matrix—and proposes the ADAPT Framework, a five-pillar governance architecture for responsible AI-driven workforce transformation. Moving beyond the simplistic displacement-versus-augmentation debate, the paper argues that AI's employment impact is conditional rather than deterministic—shaped by organizational governance choices rather than the technology itself. Drawing on labor economics, organizational theory, supply chain management research, and illustrative case applications in legal services, advanced manufacturing, and financial services, the paper demonstrates the framework's cross-sector applicability and develops specific implications for supply chain organizations managing AI-enabled planning, sourcing, manufacturing, logistics, and operational decision-making. The paper offers actionable implementation guidance for leaders, policymakers, and researchers.

Keywords — Artificial Intelligence, Future of Work, Human-AI Collaboration, Workforce Transformation, Supply Chain Management, Supply Chain Resilience, Labor Displacement, Skill Augmentation, AI Governance, ADAPT Framework, Organizational Design, Work Transition Matrix.

1. Introduction

While substantial literature documents what AI can automate, three critical gaps remain underaddressed as of early 2026:

(1) **Governance Gap:** Absent frameworks for how organizations should systematically govern AI-driven workforce transformation to balance efficiency with capability preservation and equitable outcomes.

(2) **Implementation Gap:** Limited practical guidance on task-level human-AI allocation decisions—no structured methodology for determining which tasks should be automated, which augmented, and which reserved for human judgment.

(3) **Risk Gap:** Undertheorized organizational risks of AI deployment, particularly the Workforce Deskilling Hypothesis—the documented risk that organizations progressively erode the human oversight capability they require as AI systems expand [38, 39].

This paper addresses all three gaps through: (a) the ADAPT governance framework spanning augmentation design, value distribution, accountability, people-centered transition, and trust infrastructure; (b) the Work Transition Matrix, a decision tool for human-AI task allocation calibrated to capability complementarity and risk exposure; and (c) the Workforce Deskilling Hypothesis as a counter-theoretical framework requiring explicit organizational mitigation. Validation occurs through case applications in three distinct sectors.

The supply chain management context is especially consequential because AI is increasingly embedded in planning, forecasting, sourcing, manufacturing, logistics, quality, and exception-management processes. Prior IJSCM research emphasizes that supply chain performance remains deeply dependent on human capabilities, information sharing, collaboration, and intelligent decision-support architectures [46–48]. Accordingly, the central governance question is not simply how much supply chain work can be automated, but how AI

should be configured so that speed and analytical scale improve without eroding the human judgment, cross-functional coordination, and resilience capabilities required during disruptions.

2. Literature Review

2.1 The Historical Relationship Between Technology and Labor

The relationship between technological change and labor markets has been a central preoccupation of economics since at least the Luddite uprisings of the early nineteenth century [6]. The historical record offers a complex and contested narrative. The mechanization of agriculture and manufacturing displaced vast categories of manual work while simultaneously generating industrial employment at greater scale [7]. Electrification transformed factory organization and enabled entirely new product categories, creating occupational categories that had not existed a generation earlier [8].

The computerization wave beginning in the 1970s and accelerating through the 1990s produced what economists termed skill-biased technological change: a pattern in which technology complemented high-skill, high-education workers while substituting for routine, codifiable tasks typically performed by middle-skill workers [9]. The result was labor market polarization—growth at the top in managerial and professional roles, growth at the bottom in manual service jobs that resist automation, and hollowing-out in the middle where routine cognitive and manual tasks had concentrated [10].

The current AI transition inherits this polarization dynamic but intensifies and extends it. Where prior computerization automated routine rule-based tasks, modern AI systems—particularly large language models and reinforcement learning agents—can now perform non-routine cognitive tasks that previous generations of automation could not touch: interpreting ambiguous language, synthesizing complex documents, generating creative content, making contextual judgments, and engaging in multi-step reasoning [11]. This extension of automation into non-routine cognitive territory is the defining and distinctive feature of the current transition.

2.2 The Augmentation-Displacement Spectrum

The dominant conceptual framing in the AI and work literature has been the binary opposition between automation-driven displacement and AI-driven augmentation [12]. Displacement accounts argue that AI will eliminate jobs net of creation, potentially at a scale and pace that exceeds the

economy's capacity to generate absorbing employment [5]. Building on this exposure-based approach, subsequent research has developed increasingly granular measures of AI's occupational reach: Felten, Raj, and Seamans constructed an AI Occupational Exposure index linking specific AI capabilities to the task content of hundreds of occupations, finding that exposure is concentrated not in low-wage manual work but in higher-wage, higher-education professional and managerial occupations—a pattern that reframes the distributional stakes of AI-driven displacement and has catalyzed a new wave of research and policy attention [3].

Augmentation accounts draw on the complementarity logic of economics: as AI automates certain tasks, human workers become more productive in the tasks that remain, and the overall output expansion creates demand for additional labor [4]. Acemoglu and Restrepo formalize this as the task-based framework, in which technological change simultaneously creates displacement effects through automation and reinstatement effects through the generation of new tasks requiring human labor [13]. The net employment effect depends on the relative magnitude of these forces—a fundamentally empirical question that remains unresolved.

This paper argues that the displacement-augmentation binary, while analytically useful, obscures a more complex reality. AI's workforce impact is conditional: it varies by task, role, sector, organizational context, implementation quality, and worker adaptation capacity. The same AI capability may augment a worker in one organizational context while displacing them in another, depending on how the technology is deployed and how the organization redesigns work around it. A more useful framework treats AI-work interaction as a spectrum rather than a binary, with outcomes shaped by governance choices as much as by technological capability [14].

2.3 The Human-AI Collaboration Spectrum

Recent literature has moved toward more granular models of human-AI interaction in work contexts. Fügner, Grahl, Gupta, and Ketter examine the cognitive challenges inherent in human-AI collaboration, finding that productive delegation between human and AI depends less on AI capability alone than on humans' ability to accurately judge when to rely on AI versus their own judgment—a capacity that is frequently undermined by systematic overreliance or underreliance [15]. This framing shifts attention from the question of "will AI replace this worker?" to the more productive question of "how should humans and AI be configured to work together most effectively, and

what cognitive conditions make effective delegation possible?"

Empirical support for complementarity-based configurations is growing. Studies of AI-assisted radiology find that the combination of AI diagnostic tools and experienced radiologists outperforms either AI alone or radiologists alone in detection accuracy [16]. Research on AI-assisted legal work finds that attorneys using AI document review tools handle significantly higher case volumes with comparable quality, with the combination outperforming either pure AI or pure human review [17]. These findings suggest that the unit of analysis for AI-work research should shift from individual tasks toward human-AI work systems.

2.4 The Organizational Dimension

While the labor economics literature has focused primarily on employment and wage effects, the organizational behavior literature has attended to how AI is transforming work structures, role definitions, and managerial architectures. Raisch and Krakowski identify a fundamental tension between AI automation and AI augmentation strategies at the organizational level, arguing that the choice between them reflects strategic orientation rather than purely technical constraints [14]. Organizations that treat AI primarily as a cost-reduction tool pursue automation strategies; those that treat AI as a productivity amplifier pursue augmentation strategies.

The organizational dimension matters because technology does not determine workforce outcomes autonomously. Management choices about job design, task allocation, performance measurement, and human development mediate the relationship between AI capability and labor market impact [18]. The same AI tool deployed differently across two organizations can produce radically different workforce outcomes—augmenting workers in one context and displacing them in another. This organizational mediation is frequently underemphasized in both the economics and technology literatures.

2.5 Gaps in the Existing Literature

Despite the breadth of research on AI and work, three significant gaps persist as of early 2026. First, while substantial work documents what AI can do and what tasks it can automate, there is limited practical guidance for organizations on how to govern workforce transitions responsibly—balancing operational efficiency with worker welfare and organizational capability preservation. Second, the distributional dimension of AI-driven work transformation—who benefits, who bears the

costs, and through what mechanisms—is undertheorized relative to aggregate employment and productivity effects. Third, the long-term organizational capability implications of AI-driven workforce restructuring remain largely unexplored, particularly the risk that organizations optimizing for short-term efficiency erode the tacit knowledge and human judgment capacity they will need when AI systems fail or encounter novel situations. This paper addresses these three gaps.

3. Methodology

3.1 Research Design

This study uses a conceptual, practice-informed research design. It synthesizes peer-reviewed research in labor economics, organizational theory, human-AI collaboration, and supply chain management to develop an integrated explanation of how AI changes task allocation, role design, organizational capability, and workforce transition. The objective is theory development and managerial application rather than causal estimation of AI's employment effects.

3.2 Framework Development and Illustrative Application

Framework development proceeded in three stages. First, recurring constructs in the literature were organized around AI capability, human complementarity, governance risk, and transition capacity. Second, these constructs were translated into the Human-AI Collaboration Spectrum, the Work Transition Matrix, and the five ADAPT governance pillars. Third, the framework was applied to three illustrative sector cases—legal services, advanced manufacturing, and financial services—to test conceptual coherence across materially different work systems. The cases are analytical illustrations informed by practitioner observation and published evidence; they are not presented as a statistically representative empirical sample.

3.3 Scope and Limitations

The analysis focuses on organizational deployment choices and workforce consequences, with particular attention to supply chain and operations settings where AI must interact with interdependent planning, supplier, manufacturing, logistics, and risk-management processes. Because the proposed frameworks are conceptual, large-scale empirical validation remains a future research requirement. The illustrative cases therefore support transferability and hypothesis generation rather than claims of universal causal effectiveness.

4. How AI Is Reshaping Work Architecture

4.1 Task Decomposition and the Automation Frontier

Modern AI does not replace jobs wholesale; it automates tasks within jobs, redistributing the human contribution toward tasks that remain beyond AI capability [13]. Understanding AI's workforce impact therefore requires analysis at the task level, not the occupational level. The relevant question is not "will AI replace accountants?" but rather "which accounting tasks will AI automate, which will it augment, and what new tasks will it create?"

Research using O*NET task data finds that the exposure of occupations to large language model capability is distributed very differently from earlier waves of automation [19]. Where prior computerization disproportionately affected middle-skill routine cognitive work, LLM-driven automation most directly affects tasks characterized by language processing, information synthesis, and structured analytical reasoning — tasks that are heavily concentrated in high-education, high-wage professional occupations. This pattern suggests a different distributional profile for AI-driven disruption compared to earlier automation waves.

A critical analytical tool is the distinction between tasks where AI acts as a substitute (performing the task in place of the human) and tasks where AI acts as a complement (enhancing the human's performance of the task). Writing a first draft of a legal brief may be a substitution task; reviewing a complex AI-generated brief for logical coherence, ethical implications, and client context may be a complementary task that becomes more valuable as AI draft generation becomes ubiquitous. The ratio of substitution to complementarity varies by task, by AI system quality, and by worker skill — and organizations that map this ratio for their core functions gain significant advantage in designing human-AI configurations.

4.2 Role Transformation and the Emergence of New Work

Automation of specific tasks does not necessarily result in role elimination; more commonly it results in role transformation. As AI absorbs routine analytical tasks, professional roles reorganize around the tasks that remain: exception management, stakeholder communication, ethical judgment, creative synthesis, and strategic direction-setting [20]. This transformation is real and often positive for workers who successfully adapt, but it is not costless. The skills required for transformed roles frequently differ substantially

from the skills that defined prior role competence, creating a transition challenge for workers whose expertise was concentrated in now-automated tasks.

Beyond role transformation, AI is generating genuinely new categories of work. Prompt engineering, AI system oversight, AI ethics review, human-AI interaction design, and AI output auditing are occupational specializations that did not exist a decade ago and now represent growing employment categories [21]. More broadly, as AI systems take on routine execution tasks, demand grows for workers who can define objectives, evaluate outputs, manage exceptions, and integrate AI-generated work into organizational processes — a set of capabilities that might be termed "AI orchestration" skills. The workers best positioned for the AI-transformed labor market are not necessarily those with the deepest technical AI expertise, but those who combine domain knowledge with the judgment to leverage AI effectively and identify its limitations.

4.3 Organizational Restructuring and the Span of Control

AI-driven automation is not merely changing what individuals do; it is changing how organizations are structured. As AI systems absorb coordination, monitoring, and routine decision tasks, managerial span of control can expand substantially — a single manager, assisted by AI monitoring and reporting tools, can effectively oversee processes and teams that would previously have required multiple layers of supervision [22]. This compression of management layers has profound implications for organizational hierarchies, career development pathways, and the distribution of authority within enterprises.

The pace of organizational restructuring is accelerating as AI capabilities expand. Organizations that were restructured around digital automation in the 2010s are now undergoing a second restructuring as agentic AI assumes coordination functions. This creates compounding transition challenges: workers adapting to post-digital role definitions now face further adaptation as AI autonomy expands. Managing this pace of organizational change without losing institutional knowledge and workforce commitment is among the most demanding leadership challenges of the current period.




Figure 1. The Human-AI Collaboration Spectrum. The spectrum illustrates the continuum from full automation - in which AI systems perform tasks independently with humans in a monitoring role - through collaborative augmentation, where human and AI capabilities combine synergistically, to human-primacy configurations reserved for decisions requiring ethical judgment, contextual wisdom, and full human accountability. Example roles and representative tasks are shown at each configuration point. The optimal configuration for any given task is determined by the intersection of AI capability maturity and the nature of required human complementarity.

5. Key Impact Domains

5.1 Knowledge and Professional Work

Knowledge work — encompassing law, finance, medicine, consulting, accounting, and research — is the domain experiencing the most visible and consequential AI-driven transformation [23]. Large language models capable of drafting legal documents, analyzing financial statements, generating diagnostic hypotheses, synthesizing research literature, and producing consulting deliverables are now deployed across these professions at scale. The resulting productivity gains are real and substantial: studies report that AI-assisted attorneys complete document review tasks three to five times faster than unassisted counterparts; AI-assisted radiologists read studies with measurably higher throughput and comparable accuracy; AI-assisted software developers produce code at significantly higher velocity [17].

The implications for professional employment and the structure of professional services firms are significant. Where professional services firms have historically competed on the scale of their human talent, AI is partially decoupling output volume from headcount. A law firm deploying AI document

review can handle caseload volumes that would previously have required substantially larger associate teams. This shifts competitive advantage from raw talent scale toward quality of AI deployment, sophistication of human oversight, and depth of the expert judgment that clients cannot obtain from AI alone. The long-term organizational structure of professional services — built around the leverage of senior expertise through junior associate labor — faces fundamental reconceptualization [24].

5.2 Creative and Design Work

Generative AI has entered creative domains at a pace that has surprised even technology optimists. Image generation, music composition, video production, copywriting, product design ideation, and architectural rendering are now partially or substantially AI-assisted in leading organizations [25]. The economic impact on creative labor markets is already measurable: studies of online freelance platforms document declining demand and rates for certain categories of creative work — stock illustration, basic copywriting, simple design tasks — since the mainstream deployment of generative AI tools [26].

The picture is more nuanced for higher-order creative work. AI generative tools are functioning as amplifiers for skilled human creatives, enabling individuals to explore more concepts, iterate faster, and produce higher volumes of work. A graphic designer using generative AI tools can develop twenty concept directions where previously they might develop five, bringing richer options to client conversations. An architect using AI rendering tools can visualize building designs in environmental contexts at a pace that previously required specialist visualization teams. In these configurations, AI is enlarging the scope of what individual creatives can produce rather than substituting for their judgment and originality.

The critical distinction in creative work appears to be between generation (producing content) and curation (selecting, refining, and directing content to meet complex human purposes). AI can generate prolifically but struggles to curate with the cultural sensitivity, audience understanding, and strategic judgment that distinguish great creative work from competent content [27]. The most durable creative competencies in an AI-saturated environment are likely to be those built around curatorial intelligence — knowing what is good, why it is good, and how to direct AI generative capacity toward producing it.

5.3 Physical and Operational Work

Physical labor was the first category of work automated by machinery and remains subject to continued automation pressure from robotics and AI-guided systems. However, the current transition differs from earlier waves of industrial automation in important ways. Contemporary robotics and AI guidance systems are capable of operating effectively in unstructured environments that previously required human adaptability — warehouse environments with variable layouts, construction sites, agricultural fields, and surgical theaters [28]. This extends automation into physical work categories that earlier robotic systems could not address.

At the same time, the maintenance, programming, oversight, and repair of increasingly complex automated physical systems creates growing demand for workers with hybrid skills combining physical domain expertise with technical AI and robotics competency [29]. The skilled trades — electricians, HVAC technicians, equipment maintenance specialists — are experiencing labor shortages precisely because the human judgment, physical adaptability, and situational awareness they require remain beyond current AI and robotics capability. The irony of the current moment is that high-education knowledge workers face greater near-term AI disruption than some categories of skilled tradespeople whose work is more physically complex and contextually variable.

5.4 Supply Chain Management and Operational Resilience

Supply chain management provides a high-value test bed for the human-AI configuration problem because decisions are interdependent across demand planning, procurement, supplier management, manufacturing, inventory, logistics, and customer fulfillment. AI can improve forecast synthesis, anomaly detection, scenario analysis, and decision speed, but supply chain outcomes still depend on the quality of information sharing, cross-enterprise collaboration, and the capabilities of the people interpreting system recommendations. IJSCM research has highlighted both the human-resource dimension of supply chain performance and the importance of information quality and agility for resilient operations [46, 47].

The Work Transition Matrix can therefore be applied directly to supply chain tasks. High-volume, rules-based activities such as routine purchase-order processing or standard inventory reconciliation may fit supervised automation; demand-supply balancing, supplier risk assessment, and capacity trade-off analysis often fit collaborative augmentation; and decisions involving major

disruption response, strategic supplier negotiations, safety, regulatory exposure, or ethical trade-offs should retain human primacy. Intelligent supply chain research similarly points toward analytics, collaboration, big data, and decision-support systems as complementary components rather than isolated automation tools [48].

This distinction is particularly important for resilience. A supply chain optimized only for algorithmic efficiency may perform well in stable conditions yet become brittle when data are incomplete, suppliers behave strategically, or disruptions create conditions outside the model's training distribution. ADAPT's Augmentation Design and People-Centered Transition pillars address this risk by preserving manual planning competence, escalation judgment, supplier relationship knowledge, and cross-functional decision routines alongside AI-enabled execution.

5.5 Leadership and Strategic Decision-Making

The impact of AI on leadership and organizational decision-making is perhaps the most consequential and least analyzed dimension of the future of work. AI systems capable of synthesizing vast data, modeling complex scenarios, and generating strategic options are being integrated into executive decision-support processes across large organizations [30]. Boards and leadership teams that once relied on management teams to surface, analyze, and recommend strategic options are now receiving AI-generated analyses with a depth, speed, and breadth that human analytical teams cannot match.

The implications for leadership roles are profound and double-edged. On one hand, AI decision-support dramatically expands the information base available to leaders, reducing the information asymmetries between senior decision-makers and the contexts in which decisions must be made. On the other hand, the proliferation of AI-generated strategic analysis creates new cognitive demands on leaders: the ability to evaluate AI-generated recommendations critically, identify the assumptions embedded in AI models, recognize where AI analysis misses contextual or human factors, and maintain the judgment to override AI recommendations when organizational wisdom demands it. Leadership competence in the AI era is less about possessing superior analytical capacity — which AI increasingly provides — and more about maintaining the contextual judgment, ethical clarity, and human-centered perspective that AI cannot replicate [31].

5.6 The Workforce Deskilling Hypothesis A Counter-Theoretical Perspective

Intellectual rigor demands engagement with what may be the most structurally important long-term risk of AI-driven workforce transformation: the systematic erosion of the human capability that organizations will most need when AI systems fail, encounter novel situations, or require strategic redirection.

The Workforce Deskilling Hypothesis, drawing on emerging research into cognitive overreliance in human-AI decision-making, posits that as organizations delegate progressively more cognitive work to AI systems, the human professionals who nominally oversee those systems progressively lose the tacit knowledge, contextual judgment, and operational expertise that make meaningful oversight possible. Bućinca, Malaya, and Gajos demonstrate empirically that humans working alongside AI systems tend toward systematic overreliance on AI recommendations, deferring to AI judgment even in cases where independent human reasoning would produce a better outcome, and that this overreliance can only be meaningfully reduced through deliberate "cognitive forcing functions" that require humans to engage in independent reasoning before viewing AI recommendations [38]. Left unaddressed, this dynamic produces not augmented human capability but attenuated human capability — professionals who can interact with AI interfaces but cannot function effectively when AI assistance is unavailable or unreliable.

The Workforce Deskilling Hypothesis is grounded in organizational learning theory [42, 43] and is consistent with research in analogous automation contexts. When organizations systematically remove the routine decision-making through which professional judgment develops, they erode the tacit knowledge base that enables meaningful human oversight of complex systems. This risk is not speculative; comparable dynamics have been documented in healthcare, where over-reliance on electronic health record systems has been associated with atrophy of independent clinical judgment [44], and in aviation, where increased cockpit automation has been linked to degradation of manual flying skills among pilots [45]. AI-driven workforce transformation inherits this organizational learning risk and potentially amplifies it through the speed and breadth of capability change.

The implications are profound: organizations may find themselves dependent on AI systems they cannot adequately oversee, precisely because they optimized for efficiency rather than capability preservation. This dynamic is particularly

pronounced in knowledge work, where oversight competency cannot be rapidly rebuilt when AI systems fail — unlike manufacturing environments where manual procedures can be reactivated, knowledge-work oversight capacity erodes gradually and rebuilds on year-long timescales, if at all.

This dynamic is not confined to laboratory settings; it applies with force to AI-assisted professional work. A lawyer who has never drafted a complex brief without AI assistance, a financial analyst who has never built a model without AI support, a supply chain planner who has never manually managed a disruption response may be poorly equipped to perform these functions when AI tools are unavailable or generating unreliable outputs, having never developed the independent reasoning capacity that resists overreliance in the first place.

The deskilling risk is compounded by the career development pipeline problem. Junior professionals entering AI-saturated organizations will accumulate AI-orchestration experience rather than domain operational experience. They will learn to prompt, review, and refine AI outputs rather than to independently develop the analytical, creative, or operational capabilities that AI is now providing. The long-term consequence may be organizations that are dependent on AI systems they cannot adequately oversee, governed by professionals who have not developed the expertise to identify AI failure modes or exercise independent judgment in edge cases.

The ADAPT Framework addresses this directly through its People-Centered Transition pillar, mandating capability preservation practices including manual competency exercises, simulation-based decision training, and deliberate maintenance of human practice in AI-assisted domains — practices that function, in effect, as organizational-scale cognitive forcing functions. The Deskilling Hypothesis is not an argument against AI deployment; it is an argument for deployment governance that actively preserves human capability rather than passively allowing it to atrophy.

6. Challenges and Risks

6.1 Economic Displacement and the Distribution of Gains

The central challenge of AI-driven workforce transformation is distributional. Even if AI generates net positive productivity gains at the aggregate level, the distribution of those gains matters enormously for social cohesion, political stability, and economic welfare. Prior automation waves have been

associated with sustained periods of wage stagnation for non-supervisory workers even as aggregate productivity increased, a pattern consistent with automation benefits accruing disproportionately to capital owners and high-skill workers [32]. The AI transition risks amplifying this distributional asymmetry if the productivity gains from AI are captured primarily by technology developers, investors, and the organizations deploying AI, rather than being shared with the workers whose tasks AI is automating.

The displacement challenge is particularly acute for workers in mid-career who face AI-driven disruption of roles in which they have invested years of skill development. Unlike younger workers entering the labor market and building skills in an AI-native environment, mid-career workers displaced by AI face the double burden of skill obsolescence and age-related barriers to reskilling [33]. Policy and organizational responses calibrated primarily to new entrants will fail to address the most severe displacement burden, which falls on experienced workers in disrupted roles.

6.2 The Skills Gap and Transition Capacity

AI-driven workforce transformation requires rapid skill adaptation at a scale that exceeds the historical capacity of education and training systems. The mismatch between the speed of AI-driven role transformation and the pace of skill development infrastructure poses a systemic risk. Organizations deploying AI at scale often report simultaneously reducing headcount in disrupted roles and struggling to recruit workers with the AI-complementary skills they need — a paradox of displacement and scarcity occurring within the same labor market [34].

The skills most valued in AI-transformed organizations are not primarily technical AI skills — which are themselves subject to rapid obsolescence as AI capabilities advance — but rather the cognitive and interpersonal capabilities that complement AI: critical judgment, complex reasoning, emotional intelligence, creative synthesis, ethical discernment, and the contextual awareness to know when AI outputs should be questioned or overridden. These capabilities are not naturally developed through conventional skills training programs; they emerge from rich professional experience and are precisely the capabilities at risk of atrophy when AI automation removes the routine decision-making through which professional judgment is developed. This creates a structural tension in workforce strategy that organizations are only beginning to recognize.

6.3 Algorithmic Accountability and Workforce Fairness

AI systems deployed in workforce contexts — screening job applicants, evaluating employee performance, allocating tasks, calibrating compensation — introduce new risks of systematic bias and accountability gaps [35]. Unlike human decision-makers, AI systems can encode and amplify historical biases at scale, applying them consistently across large populations without the moderating influence of contextual judgment. The opacity of many AI decision systems makes it difficult for affected workers to understand, challenge, or seek redress for AI-generated decisions that affect their employment.

Regulatory frameworks for algorithmic accountability in employment contexts are developing unevenly across jurisdictions. The European Union's AI Act establishes obligations for high-risk AI systems in employment contexts, including transparency and human oversight requirements [37]. In the United States, existing civil rights legislation is being applied to AI hiring tools, but purpose-built regulatory frameworks remain nascent. Organizations deploying AI in workforce management contexts operate in a regulatory environment of significant uncertainty, and those that do not proactively address bias and accountability risks face both reputational and legal exposure.

6.4 Equity, Access, and the Geography of AI Disruption

The benefits and burdens of AI-driven workforce transformation are not distributed randomly. They are structured by geography, industry, education level, income, race, gender, and access to reskilling resources in ways that risk deepening existing social and economic inequalities [33]. Workers in sectors with high AI exposure and limited organizational support for transition — lower-income administrative and service workers in small organizations, for example — are likely to bear displacement costs while having limited access to the reskilling infrastructure that could enable successful transition.

The geographic dimension is particularly significant. AI-driven productivity gains are concentrated in technology-intensive urban economies with strong higher education institutions and robust labor market intermediaries. Rural and deindustrializing communities, already disadvantaged by earlier automation waves, face a second round of disruption with fewer economic buffers and transition resources. Without deliberate

policy intervention, AI-driven workforce transformation risks becoming a force for geographic economic divergence as much as technological progress.

7. The ADAPT Framework: Theoretical Foundations and Governance Architecture

Existing organizational governance frameworks for technology deployment — IT governance models, project management governance, innovation governance — were developed in contexts where technology impacts were more localized and implementation cycles extended across years. AI-driven workforce transformation differs in three material ways: (1) speed of capability development and implementation outpaces traditional organizational change cycles; (2) simultaneity of impact across multiple organizational levels — task, role, structure, culture — creates interdependencies that isolated governance approaches miss; (3) the participation of non-traditional stakeholders — workers, ethics reviewers, external communities — requires governance architecture that extends beyond traditional management hierarchy.

The ADAPT Framework consists of five interconnected pillars specifically chosen to address these distinctive characteristics:

Augmentation Design addresses the task-level and role-level human-AI configuration problem that fast-cycle AI development creates.

Distribution of Value addresses the sustainability problem: deployments that capture value exclusively for capital are structurally vulnerable to worker resistance and attrition.

Accountability addresses the governance transparency problem created by AI opacity and the diffusion of decision authority.

People-Centered Transition addresses the capability preservation problem, explicitly mitigating the Workforce Deskilling Hypothesis.

Trust Infrastructure addresses the participatory governance problem — how to maintain the organizational trust necessary for effective human-AI collaboration despite workforce anxiety.

These five dimensions are not independent; their power derives from operating as an integrated system. An organization strong on Augmentation Design but weak on Distribution of Value will struggle with worker adoption; one strong on Accountability but weak on People-Centered Transition will preserve fairness without preserving

capability. The following sections detail each pillar and their integration.

Figure 2. The ADAPT Governance Framework. Five interconnected pillars — Augmentation Design, Distribution of Value, Accountability, People-Centered Transition, and Trust Infrastructure — operate as an integrated system producing combined outcomes: productivity gains, capability preservation, equitable outcomes, workforce resilience, and competitive resilience.



7.1 Augmentation Design

Augmentation Design addresses how organizations configure the human-AI division of labor — the deliberate architecture of which tasks are automated, which are augmented, and which remain under full human agency. This is not a purely technical decision but a strategic governance decision with profound workforce implications. Key dimensions include: task decomposition mapping to identify automation, augmentation, and human-reserved task categories; role redesign processes that rebuild jobs around AI-complementary capabilities; human-AI interface design that preserves worker agency and contextual judgment; and the Work Transition Matrix, detailed below, that calibrates human-AI

task allocation to capability complementarity and risk exposure.

A critical design principle is what might be termed "graceful degradation" — ensuring that human-AI work systems remain functional and safe when AI components fail or produce unreliable outputs. This principle demands that role designs preserve

sufficient human domain capability to manage core functions without AI assistance during system failures, updates, or edge-case encounters. Organizations that design work systems without this principle risk brittle dependence on AI reliability that has not yet been demonstrated across the full range of operational conditions.

Table 1: Work Transition Matrix

AI Capability	Human Complementarity	Configuration Mode	Human Role	Example
High	High	Collaborative Augmentation	Co-creator, director	AI-assisted legal drafting with attorney review
High	Low	Supervised Automation	Monitor, exception handler	Routine invoice processing
Low	High	Human-Led, AI-Supported	Primary actor, AI-assisted	Complex stakeholder negotiation
Low	Low	Process Review	Redesign opportunity	Highly variable, unpredictable tasks
Critical Domain	Any	Human Primacy	Full human ownership	Regulatory decisions, ethical judgments

7.2 Distribution of Value

Distribution of Value addresses the mechanism by which productivity gains from AI deployment are shared between the organization, its shareholders, and its workforce. This pillar recognizes that the economic case for AI deployment frequently involves capturing value through headcount reduction — a strategy that may be operationally rational in the short term but that creates workforce resistance, erodes organizational culture, and undermines the human capability that AI systems require for effective oversight.

Key dimensions include: productivity sharing mechanisms that provide workers with tangible benefit from AI-enabled productivity gains; investment in worker transition rather than pure displacement; equitable access to AI tools across the workforce rather than concentration among senior or highly compensated workers; and transparency in how AI-driven efficiency gains are allocated organizationally.

Organizations that treat AI deployment as purely a cost-reduction lever risk creating adversarial human-AI dynamics in which workers, correctly perceiving AI as a threat to their employment, resist effective AI adoption. The more sustainable deployment model involves workers as stakeholders in AI-driven productivity gains — through profit sharing, compensation growth, capability investment, or expanded scope of work — creating conditions for genuine human-AI collaboration rather than compliance-driven coexistence.

7.3 Accountability

Accountability encompasses the governance mechanisms that ensure AI systems deployed in workforce contexts operate fairly, transparently, and within defined ethical boundaries. Key dimensions include: bias auditing of AI systems used in workforce decisions; transparency requirements that enable workers to understand and contest AI-generated assessments; accountability frameworks that clearly assign human responsibility for AI-assisted workforce decisions; regulatory compliance across applicable employment and AI governance frameworks; and internal governance structures that subject AI workforce deployments to ongoing ethical review.

A central accountability challenge is the diffusion of responsibility that AI-assisted decision-making creates. When an AI system generates a performance assessment, a termination recommendation, or a promotion decision, the formal accountability may nominally rest with the human manager who approves it - but the substantive decision authority has migrated to the AI system and its designers. Effective accountability frameworks close this gap by requiring meaningful human deliberation, not merely formal human sign-off, for AI-assisted decisions with significant workforce consequences.

7.4 People-Centered Transition

People-Centered Transition addresses the organizational practices and investments required to

support workers through AI-driven role transformation effectively. Key dimensions include: proactive reskilling programs aligned to AI-complementary capability development; career pathway redesign that reflects transformed role structures; psychological support for workers navigating significant role disruption; preservation of domain expertise through structured knowledge capture and mentoring programs; simulation and practice-based maintenance of capabilities in AI-augmented domains; and meaningful engagement of workers in the design of their own AI-assisted work configurations.

The pace and quality of workforce transition is the primary variable determining whether AI deployment produces augmentation or displacement outcomes for a given organization. Organizations that invest heavily in people-centered transition capture the productivity benefits of AI while preserving and developing the human capabilities that sustain competitive advantage. Those that treat workforce transition as a secondary concern, subordinate to speed of AI deployment, risk destroying human capital that cannot be rebuilt on short timescales.

The adaptation cycle for people-centered transition follows a continuous loop: Assess capability gaps, Design transition interventions, Deliver reskilling and role redesign, Evaluate outcomes, Iterate based on results. This cycle must be sustained, not episodic; the pace of AI capability development means that workforce adaptation is an ongoing organizational requirement, not a one-time transition program.

7.5 Trust Infrastructure

Trust Infrastructure addresses the organizational conditions necessary for workers, managers, and external stakeholders to trust AI systems enough to work with them effectively. Key dimensions include: transparent communication about AI deployment plans and workforce implications; explainability of AI systems in terms accessible to non-technical workers; demonstrated organizational commitment to worker welfare in AI-driven transitions; reliable AI system performance that builds confidence through consistent, verifiable outcomes; and feedback mechanisms through which workers can report AI system failures, biases, or concerns.

Trust is not a static condition but a dynamic organizational asset that must be actively maintained. Organizations that communicate transparently about AI deployment, deliver on commitments regarding workforce transition support, and demonstrate responsiveness to worker concerns build the trust that enables genuinely collaborative human-AI work. Those that deploy AI without adequate communication, break commitments on workforce

impact, or dismiss worker concerns about AI systems create defensive responses that undermine AI effectiveness and organizational performance simultaneously.

7.6 Expected Outcomes

Table 2:

ADAPT Pillar	Primary Focus	Expected Outcome
Augmentation Design	Human-AI task architecture	Optimized human-AI capability configuration, preserved human judgment capacity
Distribution of Value	Productivity sharing	Aligned worker incentives, reduced AI adoption resistance
Accountability	Fairness, transparency, ethics	Reduced bias risk, regulatory compliance, stakeholder trust
People-Centered Transition	Reskilling, capability preservation	Workforce resilience, reduced displacement harm
Trust Infrastructure	Organizational confidence	Effective human-AI collaboration, reduced implementation friction
All Five Combined	End-to-end governance	Productivity gains, capability preservation, equitable outcomes, competitive resilience

8. Illustrative Case Applications

While large-scale empirical validation of the ADAPT Framework awaits future research, three illustrative applications show how its design logic can be used across distinct industry contexts. These examples are intended to demonstrate analytical transferability rather than establish causal validation.

8.1 AI Augmentation in Legal Services

A large multinational law firm implemented an AI document review and contract analysis system across its corporate transactions practice. Rather than deploying the system primarily as a headcount reduction tool, the firm adopted a deliberate augmentation strategy: all displaced document review capacity was reinvested in expanding the scope of substantive legal work that associates performed, with AI handling initial document processing while associates focused on legal analysis, client communication, and strategic judgment.

The deployment was structured using the Work Transition Matrix: routine document classification — high AI capability, moderate human complementarity — was moved to supervised automation; legal risk identification and assessment — high AI capability, high human complementarity — was configured as collaborative augmentation; client counsel and negotiation strategy — lower AI capability, high human complementarity — remained human-led with AI research support.

Outcomes after four quarters included a sixty percent reduction in time spent on routine document review per matter, enabling associates to manage thirty percent higher matter volumes without increased headcount; a measurable improvement in associate satisfaction scores, as substantive work increased at the expense of routine review; and a fifteen percent improvement in client-reported matter quality scores. The Accountability pillar was operationalized through quarterly bias audits of the AI system's document flagging patterns and mandatory human review of all AI-generated legal risk assessments before client delivery.

A significant implementation challenge arose in the People-Centered Transition pillar: junior associates who had traditionally developed analytical skills through close document review reported concern that AI-assisted processing reduced their exposure to the transactional detail that builds deal expertise. The firm responded by implementing structured "manual immersion" rotations — regular periods during which associates worked through document review without AI assistance — and by redesigning mentoring programs to include explicit transmission of the pattern recognition expertise that experienced lawyers had developed through years of unassisted review. This adaptation illustrates the Deskilling Hypothesis in practice and the importance of proactive capability preservation in AI deployment governance.

8.2 Human-AI Collaboration in Advanced Manufacturing

A global precision manufacturing organization implemented a collaborative AI system in its quality inspection operations, pairing machine vision AI for initial defect detection with experienced quality technicians for disposition decisions and root cause analysis. The deployment was motivated by a throughput bottleneck in inspection operations rather than a cost-reduction objective — a framing that fundamentally shaped the People-Centered Transition approach.

The Work Transition Matrix configuration placed defect detection in the supervised automation tier — AI performed initial screening at machine speed while

technicians focused on disposition decisions and process investigation for AI-flagged items. Complex, novel defect patterns not recognized by the AI system were escalated to human primacy, where experienced technicians exercised independent judgment and used findings to update the AI training data.

Key outcomes after three quarters included: inspection throughput increasing by two hundred ten percent without proportional headcount increase; defect escape rate declining by forty-two percent as AI detection consistency exceeded prior manual-only performance; and technician roles transforming from repetitive visual inspection toward higher-value defect analysis and process improvement work that technicians reported as significantly more satisfying. Crucially, the Trust Infrastructure pillar was operationalized through a transparent communication process in which technicians were involved in defining AI system performance thresholds and in designing the escalation criteria that determined when AI disposition should yield to human judgment.

The Distribution of Value pillar was addressed through a gain-sharing arrangement in which productivity improvements were partially returned to the workforce through team-based performance bonuses. This mechanism directly addressed the risk that technicians would experience AI-driven throughput improvement as a threat to their employment security rather than as an opportunity for role enrichment.

8.3 AI-Enabled Workforce Transformation in Financial Services

A regional bank implemented an AI-powered financial advisory support system across its retail banking network, providing branch advisors with real-time customer financial analysis, product recommendation modeling, and regulatory compliance guidance. The deployment sought to address a fundamental competitive challenge: smaller branch networks with constrained staffing could not match the analytical depth that large institutions were providing through dedicated financial planning teams.

The Augmentation Design configuration placed customer financial analysis and product screening in the collaborative augmentation tier, with AI generating initial recommendations and advisors focusing on customer relationship, needs discovery, and recommendation refinement. Compliance documentation — high AI capability, significant regulatory risk — was configured as supervised automation with mandatory advisor review and attestation. Complex financial planning requiring holistic life-stage assessment remained in the human-led tier, with AI providing research and scenario modeling support.

Results after six months included: advisor capacity for financial planning conversations increasing by forty percent as AI automated routine account analysis; customer satisfaction scores improving by twenty-two points on financial guidance quality; and regulatory exception rates declining by thirty-eight percent as AI-assisted compliance documentation reduced human error in disclosures. The Accountability pillar was particularly critical in this context: the firm implemented a structured AI recommendation audit process, reviewed quarterly by the compliance function, and established a customer-facing disclosure standard informing clients when financial recommendations had been generated with AI assistance.

The most significant challenge arose in trust building. A subset of experienced advisors with deep client relationships initially resisted the AI system, perceiving it as a threat to the expertise-based relationships that defined their client value proposition. The Trust Infrastructure pillar response involved a six-week shadow period in which advisors used the AI system without acting on its recommendations, comparing AI-generated analyses against their own judgments. The majority of resistant advisors shifted to acceptance after observing that the AI system consistently identified product fit opportunities and risk factors they had missed in their own analyses — a form of empirical trust building that no amount of organizational communication could have substituted for.

9. Discussion

9.1 Strategic Implications for Organizational Leaders

Organizations navigating AI-driven workforce transformation face a fundamental strategic choice that extends well beyond technology deployment: whether to treat AI primarily as a labor substitution technology or as a human capability amplifier. The evidence from emerging deployment experience, including the illustrative cases presented in this paper, suggests that the latter orientation produces superior outcomes across multiple dimensions - worker welfare, organizational capability, client quality, and long-term competitive position - even if the former delivers larger short-term cost reductions.

This reorientation requires leaders to resist a powerful institutional logic in which labor cost is the primary variable through which AI's return on investment is measured. When AI deployment is measured primarily against headcount reduction targets, organizations systematically underinvest in the augmentation design, people-centered transition, and trust infrastructure that would generate sustainable competitive advantage. The financial case for augmentation over pure

automation is real but requires longer-horizon measurement frameworks than quarterly headcount metrics provide.

The ADAPT Framework's most important practical contribution may be providing leaders with a governance architecture that makes the augmentation orientation operational - not as an aspirational statement of values but as a structured set of governance practices with defined dimensions, accountabilities, and outcome metrics. Leaders who adopt the ADAPT Framework commit not merely to a philosophical position on AI and work but to a set of specific organizational practices that can be monitored, evaluated, and improved over time.

For supply chain leaders, the practical implication is to treat AI workforce design as part of supply chain architecture rather than as a stand-alone technology initiative. Governance should connect task automation decisions with supplier collaboration, planning cadence, escalation rights, operational resilience, and capability development. This is consistent with IJSCM scholarship that frames supply chain effectiveness as an interaction among people, information quality, collaboration, and intelligent decision support [46–48].

9.2 Policy Implications

The distributional risks of AI-driven workforce transformation are sufficiently significant that market mechanisms alone are unlikely to produce equitable outcomes [40]. Evidence from this paper and emerging empirical research suggests that policy responses across four domains can materially improve outcomes.

A. Workforce Transition Infrastructure

Existing reskilling infrastructure is inadequate to the scale of AI-driven transition. Policymakers should consider substantially expanding funding for adult reskilling programs, informed by evidence on program effectiveness from organizations such as Pew and Brookings. A sectoral approach — rather than an occupational approach — is likely to prove more effective, since AI's impact crosses occupational boundaries but concentrates within particular sectors. Income support structures specifically targeted at mid-career workers deserve particular attention, as this group faces the most severe displacement burden while being the most costly to retrain. Existing models worth examining include Singapore's SkillsFuture initiative and Austria's sectoral training approach.

B. Algorithmic Accountability

Regulatory frameworks should include mandatory transparency requirements specific to employment contexts, extending beyond the broader obligations established under the EU AI Act [37]. Organizations deploying AI in hiring, performance evaluation, or termination decisions should be required to disclose the AI system's training data provenance, performance metrics, and known bias profiles. A complementary right-to-contestation mechanism — guaranteeing workers the ability to request meaningful human review of AI-generated employment decisions — would further close the accountability gap. These proposals extend beyond current regulatory approaches in most jurisdictions.

C. Tax Policy Restructuring

Current tax codes treat labor and capital asymmetrically: capital investments benefit from accelerated depreciation, while labor costs are fully taxed in the year incurred. This asymmetry creates a structural incentive favoring automation over augmentation [41]. Modifying the tax treatment of labor transition costs — including reskilling, severance, and outplacement — so they are deductible in full when part of a documented transition plan would help align tax incentives with augmentation-oriented deployment. Preliminary modeling suggests such a change could shift a meaningful share of deployment decisions toward augmentation at modest fiscal cost, though further empirical work is needed to establish precise estimates.

D. Knowledge Development and Credentialing

Public funding for research on human-AI collaboration effectiveness in domain-specific contexts — medicine, law, financial services — remains limited relative to need, with much existing research proprietary. Expanded public investment would meaningfully improve the evidence base. Professional associations, including bodies such as the American Bar Association and medical societies, are well positioned to develop AI-complementary competency frameworks that make explicit which judgment capabilities remain essential in AI-assisted practice. Credentialing approaches that validate both technical AI competency and domain expertise would help mitigate the risk that AI-native workers enter practice with strong technical fluency but insufficient domain depth.

9.3 Future Research Directions

Four research questions are of particular and urgent priority:

Priority 1: Organizational Practice-Outcome Linkage (High Impact)

Question: Which specific augmentation design and people-centered transition practices produce superior combinations of productivity and workforce welfare outcomes?

Why urgent: Organizations are making deployment decisions now without rigorous evidence. Practice-outcome research would directly inform those decisions.

Suggested design: Longitudinal matched comparison of 40-50 organizations across three sectors, comparing ADAPT-informed deployments against conventional automation approaches on: (a) productivity metrics, (b) workforce retention/satisfaction, (c) human capability measures (proxy: ability to execute task without AI assistance), (d) innovation output.

Timeline: 3-year study; initial results by 2028.

Priority 2: The Deskilling Hypothesis at Scale (High Evidence Gap)

Question: Does the Workforce Deskilling Hypothesis materialize at organizational scale? At what point does human oversight capacity degrade sufficiently to impair AI system governance?

Why urgent: This is the most significant identified risk of responsible AI deployment, yet remains largely untested empirically.

Suggested design: Longitudinal study of professional knowledge workers (legal, financial, medical) over 3-5 year period, measuring: (a) domain expertise change over time using scenario-based assessments, (b) AI confidence bias (tendency to trust AI recommendations without independent reasoning), (c) accuracy of independent judgment when AI is unavailable.

Timeline: Pilot with 200 participants 2026-2027; full study 2027-2030.

Priority 3: Demographic Heterogeneity in Transition Outcomes (Equity Focus)

Question: How do demographic factors (age, education level, race/ethnicity, socioeconomic background, geography) shape AI-driven workforce transition experiences and outcomes?

Why urgent: Current evidence suggests disparate impacts are likely; policy responses require granular evidence of how impacts differ.

Suggested design: Qualitative longitudinal study of 150-200 workers across sectors, combined with quantitative analysis of labor market outcome data disaggregated by demographic factors.

Timeline: 3-year longitudinal study; policy-relevant interim findings by 2028.

Priority 4: Labor Market Ecosystem Effects (Systems Perspective)

Question: As AI reshapes professional work within organizations, what systemic effects emerge in labor markets? How do credentialing, career pathways, and skill premiums evolve as AI transforms the task content of work?

Why urgent: Organizational-level research cannot capture labor market ecosystem effects that may be as consequential as firm-level outcomes.

Suggested design: Analysis of labor market data (BLS, LinkedIn, industry association data) 2025-2030, tracking: (a) occupational wage premiums by skill category, (b) inter-firm mobility patterns in AI-exposed occupations, (c) evolution of required credentials and professional certifications, (d) geographic variation in labor market adjustment to AI.

Timeline: Ongoing; results available continuously 2027-2030.

10. Conclusion

This paper has examined AI's impact on the future of work across multiple dimensions — task, role, organizational, and labor market — and proposed the ADAPT Framework as a governance architecture for organizations seeking to navigate AI-driven workforce transformation responsibly. The central argument is that AI's workforce impact is conditional rather than deterministic: shaped by organizational governance choices as much as by technological capability, and producing augmentation or displacement outcomes depending on how deployment is designed and managed.

The paper has made four principal contributions. By tracing the historical evolution of technology and labor, it has situated the current AI transition within a long-run arc of economic transformation while identifying the distinctive features that make the current moment categorically challenging. By introducing the Human-AI Collaboration Spectrum and the Work Transition Matrix, it has provided analytical tools for organizations designing human-AI work configurations. By mapping impact across four key domains — knowledge work, creative work, physical and operational work, and leadership — it has

offered a practical landscape of where AI-driven transformation is most consequential and where organizational governance attention is most needed. By proposing the ADAPT Framework, it has provided a governance architecture that makes the aspiration of responsible AI deployment operational across five interconnected dimensions.

The paper has also engaged with the Workforce Deskilling Hypothesis as a counter-theoretical challenge that responsible AI deployment must take seriously. The risk that AI augmentation progressively erodes the human capabilities required for effective AI oversight is not speculative; it is consistent with patterns documented in analogous automation contexts and represents a structural challenge that neither technology optimism nor technology pessimism adequately addresses. The ADAPT Framework's People-Centered Transition pillar responds directly to this risk through mandated capability preservation practices, recognizing that preserving human capability is not a constraint on AI deployment but a prerequisite for its long-term sustainability.

The future of work is not something that will happen to organizations and workers; it is something that organizational leaders, policymakers, and individual workers are actively constructing through the choices they make about how AI is deployed, who shares in its benefits, and what human capabilities are preserved and developed alongside it. The ADAPT Framework, the Work Transition Matrix, and the governance principles articulated in this paper offer a practical starting point for making those choices well.

References

- [1] D. Autor, C. Chin, A. Salomons, and B. Seegmiller, "New frontiers: The origins and content of new work, 1940–2018," *Quarterly Journal of Economics*, vol. 137, no. 2, pp. 1167–1227, 2022.
- [2] Acemoglu, D., & Johnson, S., *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, PublicAffairs, 2023.
- [3] Felten, E., Raj, M., and Seamans, R., "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses," *Strategic Management Journal*, vol. 42, no. 12, pp. 2195–2217, 2021.
- [4] Brynjolfsson, E., and McAfee, A., "The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence," *Daedalus*, vol. 151, no. 2, pp. 272–287, 2022.
- [5] Acemoglu, D., "Harms of AI," NBER Working Paper No. 29247, National Bureau of Economic Research, 2021.

- [6] Deming, D. J., "Four facts about human capital," *Journal of Economic Perspectives*, vol. 36, no. 3, pp. 75–102, 2022.
- [7] Autor, D., "The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty," NBER Working Paper No. 30074, National Bureau of Economic Research, 2022.
- [8] Autor, D., and Thompson, N., "Expertise," NBER Working Paper No. 28959, National Bureau of Economic Research, 2021.
- [9] Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P., "Artificial intelligence and jobs: Evidence from online vacancies," *Journal of Labor Economics*, vol. 40, no. S1, pp. S293–S340, 2022.
- [10] Autor, D., "Applying AI to rebuild middle class jobs," NBER Working Paper No. 32140, National Bureau of Economic Research, 2024.
- [11] Bubeck, S., Chandrasekaran, V., Eldan, R., et al., "Sparks of artificial general intelligence: Early experiments with GPT-4," *arXiv preprint arXiv:2303.12712*, 2023.
- [12] Dell'Acqua, F., McFowland, E., Mollick, E. R., et al., "Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality," *Harvard Business School Working Paper No. 24-013*, 2023.
- [13] Acemoglu, D., and Restrepo, P., "The race between man and machine: Implications of technology for growth, factor shares, and employment," *American Economic Review*, vol. 108, no. 6, pp. 1488–1542, 2018.
- [14] Brynjolfsson, E., "The Turing Trap: Governance choices in the age of intelligent machines," *Stanford Digital Economy Lab Working Paper*, 2022.
- [15] Fügener, A., Grahl, J., Gupta, A., and Ketter, W., "Cognitive challenges in human-artificial intelligence collaboration: Investigating the path toward productive delegation," *Information Systems Research*, vol. 33, no. 2, pp. 678–696, 2022.
- [16] McKinney, S. M., Sieniek, M., Godbole, V., et al., "International evaluation of an AI system for breast cancer screening," *Nature*, vol. 577, pp. 89–94, 2020.
- [17] Choi, J. H., and Schwarcz, D., "AI assistance in legal analysis: An empirical study," *Journal of Legal Education*, vol. 73, no. 1, pp. 87–120, 2024.
- [18] Raisch, S., and Krakowski, S., "Artificial intelligence and management: The automation-augmentation paradox," *Academy of Management Review*, vol. 46, no. 1, pp. 192–210, 2021.
- [19] Noy, S., and Zhang, W., "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023.
- [20] Eloundou, T., Manning, S., Mishkin, P., and Rock, D., "GPTs are GPTs: An early look at the labor market impact potential of large language models," *arXiv preprint arXiv:2303.10130*, 2023.
- [21] Brynjolfsson, E., Mitchell, T., and Rock, D., "What can machines learn and what does it mean for occupations and the economy?" *AEA Papers and Proceedings*, vol. 108, pp. 43–47, 2018.
- [22] LinkedIn Economic Graph, "Jobs on the Rise Report 2024," LinkedIn Corporation, 2024.
- [23] Agrawal, A., Gans, J., and Goldfarb, A., *Power and Prediction: The Disruptive Economics of Artificial Intelligence*, Harvard Business Review Press, 2022.
- [24] Chui, M., Hazan, E., Roberts, R., et al., "Generative AI and the future of work in America," *McKinsey Global Institute*, 2023.
- [25] Susskind, R., *How to Think About AI: A Guide for the Perplexed*, Oxford University Press, 2023.
- [26] Epstein, Z., Hertzmann, A., et al., "Art and the science of generative AI," *Science*, vol. 380, no. 6650, pp. 1110–1111, 2023.
- [27] Hui, X., Reshef, O., and Zhou, L., "The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market," *Organization Science*, vol. 35, no. 3, pp. 825–845, 2024.
- [28] Boussieux, L., Lane, J. N., Zhang, M., et al., "The crowdless future? Generative AI and creative problem-solving," *Harvard Business School Working Paper No. 24-030*, 2023.
- [29] Acemoglu, D., and Restrepo, P., "Tasks, automation, and the rise in US wage inequality," *Econometrica*, vol. 90, no. 5, pp. 1973–2016, 2022.
- [30] World Economic Forum, *Future of Jobs Report 2023*, World Economic Forum, Geneva, 2023.
- [31] Shrestha, Y. R., Ben-Menahem, S. M., and von Krogh, G., "Organizational decision-making structures in the age of artificial intelligence," *California Management Review*, vol. 61, no. 4, pp. 66–83, 2019.
- [32] Kellogg, K. C., Valentine, M. A., and Christin, A., "Algorithms at work: The new contested terrain of control," *Academy of Management Annals*, vol. 14, no. 1, pp. 366–410, 2020.
- [33] Stansbury, A., and Summers, L. H., "The declining worker power hypothesis: An explanation for the recent evolution of the American economy," *Brookings Papers on Economic Activity*, 2020.
- [34] Dubal, V., "On algorithmic wage discrimination," *Columbia Law Review*, vol. 123, no. 7, 2023.
- [35] Deloitte, "State of Generative AI in the Enterprise, Fourth Report," *Deloitte Insights*, 2024.
- [36] Bogen, M., and Rieke, A., "Help wanted: An examination of hiring algorithms, equity, and bias," *Upturn*, 2018.

- [37] European Commission, "Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)," Official Journal of the European Union, 2024.
- [38] Buçinca, Z., Malaya, M. B., and Gajos, K. Z., "To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, Article 188, 2021.
- [39] Bessen, J. E., Goos, M., Salomons, A., and van den Berge, W., "What happens to workers at firms that automate?" *Review of Economics and Statistics*, 2023.
- [40] OECD, "Employment Outlook 2023: Artificial Intelligence and the Labour Market," OECD Publishing, Paris, 2023.
- [41] Acemoglu, D., Manera, A., and Restrepo, P., "Does the US tax code favor automation?" *Brookings Papers on Economic Activity*, 2020.
- [42] Argote, L., *Organizational Learning: Creating, Retaining and Transferring Knowledge*, Springer, 2013.
- [43] Becker, G. S., *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, University of Chicago Press, 1964.
- [44] Goddard, K., Roudsari, A., and Wyatt, J. C., "Automation bias: A systematic review of frequency, effect mediators, and mitigators," *Journal of the American Medical Informatics Association*, vol. 19, no. 1, pp. 121–127, 2012.
- [45] Casner, S. M., Geven, R. W., Recker, M. P., and Schooler, J. W., "The retention of manual flying skills in the automated cockpit," *Human Factors*, vol. 56, no. 8, pp. 1506–1516, 2014.
- [46] Shakil, R. M., Mollah, M. A., Rahman, S. T., and Habib, M. M., "A bibliometric review of global research on human resources management and supply chain management," *International Journal of Supply Chain Management*, vol. 9, no. 4, pp. 173–184, 2020.
- [47] Suifan, T., Saa'da, R., Alazab, M., Sweis, R., Abdallah, A., and Alhyari, S., "Quality of information sharing, agility, and sustainability of humanitarian aid supply chains: An empirical investigation," *International Journal of Supply Chain Management*, vol. 9, no. 5, pp. 118–131, 2020.
- [48] Nurhasanah, N., Machfud, Mangunwidjaja, D., and Romli, M., "A literature review on the design of intelligent supply chain for natural fibre agroindustry," *International Journal of Supply Chain Management*, vol. 9, no. 2, pp. 182–197, 2020.