

A Study of AI Application Prospects and Risks in the Age of AGI

From Expanding Capabilities to Civilizational Consequences

AI Applications, Opportunities, Risks, and Civilizational Consequences in the Age of AGI

AI's real impact is not the number of applications added, but the growing share of knowledge, judgment, organization, and action that passes through AI.

In-depth Study of Real-World Convergence V1.2

Author: Zijun Fu

Produced by: Civilization Leap Research Group

Evidence cutoff: August 26, 2026

Abstract translated from the Chinese original

Artificial intelligence is moving beyond demonstrations of individual model capabilities into real organizations, research, education, healthcare, government, culture, relationships, and security systems. The 2026 AI Index shows that AI has become a widely adopted general-purpose technological condition: business adoption, scientific research, and educational use are all expanding rapidly, although real deployment of agentic systems still lags considerably behind demonstrations of model capabilities [1]. The International Labour Organization's task-level research on generative AI indicates that approximately one-quarter of employment worldwide falls within some degree of exposure. At present, however, the better-supported judgment is that job content and forms of organization are being restructured; exposure should not simply be equated with the disappearance of jobs [2]. Together, these facts suggest that the civilizational significance of AI applications cannot be inferred directly from whether AI can perform a task. We must go on to ask how capabilities, once introduced into real tasks, change organizations, human-machine divisions of labor, resources, power, relationships, and people's place in the world.

This paper is neither a list of industry trends nor an attempt to make risk governance the goal of civilization. It begins with seven types of real-world settings: research and knowledge, labor and the economy, education and development, healthcare and life, public governance, relationships and culture, and security and war. It examines application prospects and risks in each, then traces how these changes travel through other domains to produce feedback and aggregate effects. Its central judgment is that the same AI capability may expand human understanding, creativity, care, and public-service capacity, or leave people indefinitely confined to final confirmation, minor editing, passive responses, and interfaces involving little responsibility. Applications that are individually reasonable may, in combination, produce shared infrastructure dependencies, outsourced judgment, drifting responsibility, concentrated power, correlated failures, and difficulties of exit. What needs investigation is not merely whether an application improves efficiency, but whom it ultimately strengthens, whose responsibilities it changes, what it hands over to systems, and what kind of civilization these changes jointly shape. Reality unfolds through causality, not through disciplines; disciplines pursue truth separately, while civilization bears the consequences together.

The paper also brings the sources of human meaning into application research. Work, creativity, knowledge, development, care, relationships, public participation, culture, faith, and the life course can all provide meaning. By changing how people enter these activities, the paths through which they develop, and the structures that sustain them, AI also changes the sources and distribution of structural meaning. The paper therefore identifies two real-world directions that require continuing comparison. In one, AI expands human capabilities, giving more people time, knowledge, participation, and a more substantial place in lived reality. In the other, system capabilities continue to deepen while human judgment, responsibility, development, and the joint formation of collective direction gradually diminish. Neither direction is technologically inevitable. Both depend on organizational design, institutional choices, the distribution of power, cultural norms, and deliberate human choices.

Keywords: AI applications; age of AGI; cross-disciplinary causality; systemic risk; human-AI collaboration; structural meaning; public power; Civilization Leap

Abstract

Artificial intelligence is moving from model-level capability demonstrations into real organizations, science, education, health care, government, culture, social relationships, and security systems. The key civilizational question is therefore no longer simply what AI can do, but what happens after capability enters real tasks and becomes embedded in institutions, resource allocation, authority, relationships, and human life. This paper studies seven application domains and then traces cross-domain feedback and aggregation. It argues that the same capability can augment human understanding, creativity, care, and public service, or progressively relocate humans to shallow roles of confirmation, correction, and response. Individually reasonable applications may also aggregate into common dependencies, correlated failures, diffuse responsibility, concentrated power, and reduced exit capacity. The paper further examines how AI reorganizes the structural sources through which people gain roles, responsibility, contribution, growth, belonging, and meaningful participation. It concludes with a falsifiable research agenda that asks not only whether AI increases performance, but whose capabilities become thicker, which responsibilities remain real, whether alternative paths remain open, and what kind of civilization emerges from the aggregate.

Keywords: artificial intelligence applications; AGI era; cross-domain causality; systemic risk; human-AI collaboration; structural meaning; public power; civilizational transition

Research Positioning, Scope, and Standards of Argument

The term "age of AGI" in this paper does not presuppose an already completed AGI entity with a universally agreed definition. It refers to the continuing entry of advanced AI into knowledge production, organizational action, resource allocation, and consequential decision chains, increasingly reconnecting in practice issues that different domains could previously address separately. For factual judgments about current systems, the paper uses specific terms such as "current advanced AI," "generative AI," and "agentic systems." Future AGI is treated only as an object of prospective research.

The paper distinguishes four types of claims. Category A comprises relatively well-supported facts, formal institutional materials, and verifiable application states. Category B comprises frontier research, competing explanations, and empirical findings still requiring replication. Category C comprises cross-disciplinary connections, future scenarios, and structural judgments requiring real-world validation. Category F comprises the philosophical and normative commitments for which the author openly takes responsibility. Category A does not automatically establish B, C, or F, and Category F must not masquerade as scientific fact. Judgments in this paper that stronger system capabilities may become decoupled from stronger capabilities of human subjects, that applications across domains may aggregate into civilizational dependencies, and that sources of meaning may shift structurally are chiefly Category C claims. The position that technology should expand human capabilities and

cannot automatically acquire authority over meaning or public power by virtue of superior capability is a Category F normative commitment.

Type of claim	Typical content in this paper	What it can support	What it cannot automatically establish
A Established facts / formal materials	Adoption rates, task exposure, formal regulations, published experiments, and actual application scope	Factual judgments bounded by object, time, and task	Complete social consequences, subjecthood, legitimacy, or civilizational direction
B Frontier research / competing explanations	Heterogeneous educational effects, long-term effects of AI companions, and limits of scientific creativity	Competing mechanisms and directions for further research	Universal conclusions presented as already settled
C Cross-domain structural judgments	Organizational dependence → changes in power; aggregation of applications → systemic dependence; reorganization of sources of meaning	Refutable cross-disciplinary hypotheses and stress tests	Claims disguised as proven by a single discipline
F Philosophical / normative commitments	A person's basic standing cannot be purchased by market value or model capability; public power requires separate justification	Why the paper cares about particular consequences and directions	Replacement of empirical research, legal procedures, or disciplinary facts

In specific deployments with traceable responsibility, explicit authorization, auditability, correction, and exit, current advanced AI may perform functional roles in understanding, protection, transmission, analysis, generation, and capability extension. These roles describe deployment relationships. They do not establish subjective experience, nontransferable continuity of a subject, or awareness of consequences for itself; nor do they automatically confer legal personhood, political membership, public office, or sovereignty. Future subjecthood remains an open question, while present institutional standing and public power must each be independently justified.

This paper does not ask whether AI is ultimately good or bad. It asks why the same capability, entering different structures, leads toward different civilizational outcomes.

Contents

Introduction | Why AI Is Becoming a Civilizational Question

1 | Capabilities Enter Reality: From Model Performance to Real Action Chains

2 | Research and Knowledge: Discovery Speed, Knowledge Infrastructure, and Research Direction

3 | Labor and the Economy: Productivity, Jobs, Human-Machine Divisions of Labor, Distribution, and Being Needed

4 | Education and Development: Access to Knowledge, Teachers' Roles, Judgment, and Intergenerational Continuity

5 | Healthcare and Life: Clinical Capabilities, Responsibility, Decisions About Life, and Caring Relationships

6 | Public Governance: Prediction, Decision Support, Dependence, Legitimacy, and De Facto Power

7 | Relationships, Culture, and Spiritual Life: Companionship, Translation, Creation, Remembrance, and Real Relationships

8 | Security and War: Accelerated Decisions, Escalating Competition, Miscalculation, and Shared Survival

9 | How Prospects Become Civilizational Benefits: What Must Follow Capability Gains

10 | From Sectoral Effects to Cross-Domain Feedback: Why Changes Travel Across Domains

11 | Aggregate Risks: How Multiple Local Optima Can Produce a Fragile System

12 | Reorganizing the Sources of Human Meaning: Occupation, Relationships, Development, Responsibility, and Public Participation

13 | Two Real-World Directions: Expanding Human Capabilities or Confining People to Shallow Interfaces

14 | What Evidence Could Change This Paper: An Agenda for Cross-Disciplinary Research and Real-World Validation

Conclusion | Technology Provides Capabilities; Civilization Bears the Consequences

Appendix A | Application-Discipline-Civilizational Consequence Matrix

Appendix B | Checks for External Circulation and Reverification of Changing Facts

Declaration on the Use of Generative Artificial Intelligence

References

Introduction | Why AI Is Becoming a Civilizational Question

Discussions of a technology could often begin by asking which capabilities it improved, then turn to adoption within industries. AI is changing that sequence. Once generative models, tool use, agents, multimodal systems, and automation infrastructure enter real settings, they do more than speed up a particular step. They penetrate organizational divisions of labor, allocations of responsibility and resources, educational pathways, public services, and social relationships. Stanford HAI's 2026 AI Index reports continuing growth in AI adoption among surveyed organizations and the entry of generative AI into many business functions, while agentic AI deployment remains at an early stage in most functions [1]. Two realities therefore coexist: capability expansion is already rapid, while comprehensive organizational restructuring is still taking shape. The crucial window lies precisely between the two.


What technology can do is a computer-science question. Why organizations adopt it is a question of management, economics, and institutions. Who gains resources after adoption, who loses jobs, and who controls data and infrastructure are political-economic questions. When AI participates in education, healthcare, and public decisions, law, ethics, politics, and professional responsibility return to the analysis. When work, relationships, culture, and developmental pathways are rewritten, questions of sociology, education, psychology, and meaning arise again. Real-world consequences do not stop at the boundary of any one discipline.

Accordingly, this paper proceeds from practical reality toward civilization. It starts with a concrete fact about an application, then traces how entry into an organization changes the human-machine division of labor. It follows how that division changes returns, responsibility, power, relationships, and human development. Finally, it examines whether changes occurring simultaneously across domains generate new aggregate outcomes. The civilizational scale here is not a higher court of judgment. It places facts established separately back within the same world.

The events at issue are civilizational in scale. They are not peripheral background for any single discipline, but a new reality into which disciplines are moving together. Each discipline must still establish its own facts and professional judgments independently. Yet it cannot stop once its own conclusion holds: it must continue tracing

how those judgments enter other domains, generate feedback, and produce new shared consequences in combination. Reality unfolds through causality, not through disciplines. What converge are real-world problems and causal chains, not disciplinary standards of evidence.

Figure 1. From AI applications to civilizational consequences: the paper's basic chain of observation



This chain also explains why "prospects" and "risks" cannot be separated into unrelated lists. The same automation system can reduce tedious work or eliminate the intermediate steps through which judgment is trained. The same AI teacher can extend access to high-quality education or reduce students' experience of difficulty, discussion, and relational learning. The same clinical support system can lower informational and professional barriers or amplify automation bias. The same government AI system can improve service responsiveness or make de facto power more dependent on opaque technological infrastructure. Opportunities and risks often diverge from a common source of capability.

AI's greatest promise is not merely how much it can do for people, but how much human time, capability, and possibility it can release.
AI's greatest risk is not merely what it may do wrong, but what humanity gradually hands over to it in order to obtain its capabilities.

1 | Capabilities Enter Reality: From Model Performance to Real Action Chains

1.1 Model capability is not a real-world consequence

Current advanced AI continues to improve at language, code, images, information retrieval, summarization, planning, and tool use. Yet between "a model completes a task in a test" and "a real system reliably completes a high-consequence task" lie many layers of conditions involving data, environments, permissions, tools, organizational processes, user behavior, cybersecurity, regulation, and responsibility. NIST's AI Risk Management Framework and its Generative AI Profile both situate risk management within the full lifecycle and specific contexts of use rather than treating a single model score as system trustworthiness [24][25]. Systems-engineering research has also long emphasized that complex-system accidents often arise from interactions, organizations, and control structures rather than the isolated failure of one component [30][31].

The minimum unit of real-world analysis should therefore be an action chain, not merely "a model": who frames the problem, selects the data, and invokes the model; which tools and permissions it receives; who reviews the

result and decides whether to execute it; who bears the consequences; and who can stop, repair, and exit. The more capable AI becomes, the less this chain can be compressed into "the model did it."

1.2 From assistance to dependence is a continuum, not a switch

Real organizations do not usually move overnight from "humans make every judgment" to "AI decides everything automatically." A more common progression begins with AI supplying retrieval and drafts, followed by recommendations. Once recommendations appear consistently reliable, human review gradually becomes spot-checking. Organizational processes are then redesigned around system outputs, with staffing and training changing accordingly. At that point, even if a "human in the loop" nominally remains, the actual space and time for judgment and the range of alternatives may already have contracted. Research on automation has long identified overreliance, skill degradation, supervisory burdens, and the "ironies of automation" in human-machine divisions of labor [32] [33].

Whether AI "augments people" cannot therefore be judged only by whether users complete tasks faster in the short term. We must also examine whether their capabilities increase or decline away from the system, whether they still understand the task as a whole, whether they can take over when errors occur, whether organizations retain alternatives, and whether system outputs gradually become the de facto default answer.

1.3 Four common variables in real-world applications

Variable	Initial focus	Further questions
Capability	Whether AI improves speed, quality, coverage, or access on a bounded task	Whether gains become real outcomes and remain stable across settings
Organization	How systems enter jobs, workflows, procurement, performance measures, and responsibility chains	Who gains decision authority and who retreats to end-stage interfaces
Distribution	How returns, time, costs, data, and risks are redistributed	Whether productivity gains become more widely shared capabilities for living
The subject's position	Whether people retain space for understanding, judgment, responsibility, development, and exit	Whether stronger systems also make people stronger

These four variables run through all the application settings below. They are not a new scoring system. They remind researchers that whenever "AI improves X," they must ask whose capabilities, power, responsibilities, and lives change when X enters reality.

1.4 Organizations mediate the conversion of capability into consequences

AI usually does not move directly from a model into "society." It first enters specific organizations and relational structures: businesses, schools, hospitals, governments, research institutions, platforms, and households. Organizations determine which system to procure, where to put it in a workflow, which data and tools it may access, how outcomes are measured, who holds veto power, and how saved time and costs are redistributed. The same model can produce entirely different consequences in different organizations. It can help teachers prepare lessons and diagnose learning needs, or become a substitute route through which students obtain answers. It can provide physicians with a second opinion, or become a de facto gatekeeper. It can help employees handle repetitive work, or be used to break whole jobs into numerous low-responsibility tasks.

Application research therefore cannot connect "technological capability" directly to "social outcomes." Organizational incentives and work design, at minimum, intervene. Organizations that reward only speed, cost reduction, and measurable output are more likely to use AI for substitution and compression. Those that also retain learning, responsibility, redundancy, explanation, appeals, and human development are more likely to use it for augmentation. Organizations are not neutral conduits: they mediate the direction in which capabilities are amplified.

A further Category C mechanism requires continuing organizational research: choices easiest to measure and reward in the short term may systematically favor the replacement of human subjects. Standardized AI outputs fit more readily into cost accounting, delivery-speed measures, KPIs, and compliance audits, whereas cultivating human judgment, mentoring, trial and error, redundancy, and long-term responsibility often appear first as costs in short-term budgets and performance tables. Even without intending to diminish people's substantive roles, managers facing competition, cost-cutting, accountability, and time pressures may repeatedly choose options that are more quantifiable, more standardizable, and easier to encode as procedural responsibility. If long-term outcomes support this mechanism, organizational incentives are not merely background variables in AI application; they are decisive mediators between strengthening and constraining human subjects. If organizational research instead finds that these incentives can reliably be designed in the opposite direction, this judgment should be narrowed accordingly.

This also explains why management, organizational science, and public administration cannot be absent from real-world research, even though none currently constitutes a separate paper within the sixteen in-depth disciplinary studies. Relevant facts continue to enter economics and political economy, political science, AI and computer science, sociology and anthropology, and education research within their respective domains. Reconsideration of an additional paper is warranted only if independent disciplinary facts emerge that cannot be decomposed, have no existing disciplinary owner, and are sufficient to change a civilizational judgment.

2 | Research and Knowledge: Discovery Speed, Knowledge Infrastructure, and Research Direction

2.1 Prospects: opening previously costly or complex knowledge spaces to more researchers

AI's influence on research already extends beyond literature searches and prose editing. Nature's review of scientific discovery in the age of AI notes that machine learning is being used to generate hypotheses, design experiments, process large datasets, and discover patterns difficult to obtain through conventional methods [11]. The 2026 AI Index introduces a separate Science chapter, reporting continued growth in AI-related natural-science publications across biology, chemistry, physics, and astronomy [1]. Large language models are also being discussed as supporting tools across the scientific method, from hypotheses, literature, and experimental design to interpretation of results [12].

The important prospect is not simply that scientists can write faster. A deeper change is the lowering of barriers to knowledge: small teams can access analytical, coding, simulation, and interdisciplinary retrieval capabilities previously affordable only to large institutions; researchers working in nonnative languages can reduce expression and translation costs; complex literatures can be organized faster; and candidate spaces for experiments and computations can be searched at scale. If these capabilities ultimately broaden opportunities for independent researchers and peripheral regions to enter high-quality research, AI may expand civilization's base of knowledge production rather than merely increase the output of central institutions.

2.2 Risks: accelerating knowledge does not necessarily accelerate truth-seeking

The core of research is not producing more text that resembles papers. It is asking questions worth testing, establishing reliable evidence, identifying anomalies, replicating results, and allowing failure. Existing research has not converged on whether current generative AI can independently make fundamental scientific discoveries. A 2025 Scientific Reports study found, within a specific experimental framework, that current generative AI was better at incremental discovery within existing knowledge representations, while exhibiting significant limitations in generating genuinely original hypotheses from scratch, recognizing anomalies, and making fundamental discoveries [13]. Such findings cannot be extrapolated to mean that AI will never make scientific discoveries. They do, however, show why fluent explanation, rapid generation, and genuine discovery must remain distinct.

A second risk is the concentration of research infrastructure. Frontier models, compute, data, and cloud services require enormous capital investment, and research on competition in AI infrastructure identifies substantial scale and concentration tendencies in this market [8]. If research increasingly depends on a few models, data pipelines, and cloud platforms, technological infrastructure may reshape research directions, reproducibility, access costs, and scrutiny. The democratization of knowledge and the concentration of knowledge infrastructure can occur simultaneously.

2.3 Cross-disciplinary causality: research capabilities also reshape the economy, education, and public power

Faster research accelerates discoveries in medicines, materials, energy, and industrial technologies. Discovery speed then changes industrial competition and capital allocation; industrial advantages feed back into compute and research platforms. As research tools enter education, they also change how the next generation learns science. Thus, "AI supports research" is not a benefit contained entirely within research. It may generate a feedback loop linking knowledge, capital, infrastructure, talent, and research direction.

The civilizational question in knowledge production is not whether AI can produce answers that sound more scientific, but who retains the ability to ask questions, test answers, recognize anomalies, and change research direction.

3 | Labor and the Economy: Productivity, Jobs, Human-Machine Divisions of Labor, Distribution, and Being Needed

3.1 Prospects: productivity gains may create conditions for human freedom

Generative AI has already shown measurable task-level productivity gains in some knowledge work. Noy and Zhang's experiment found improvements in speed and quality on specified writing tasks [4]. Brynjolfsson and colleagues' study of customer support found that generative AI assistance was particularly helpful to less-experienced employees [3]. Other field studies likewise reveal marked heterogeneity across tasks and capability levels [5]. These findings cannot be converted directly into macroeconomic growth, but they do indicate that, in some tasks, AI can reduce the marginal cost of knowledge work and may open capabilities previously requiring professional training to more people.

If these productivity gains shorten necessary working time, lower the cost of basic services, and strengthen small organizations and individuals, they could create genuine civilizational opportunities. More time could go to care, learning, creativity, relationships, public participation, and exploration. More people could perform complex work with lower capital requirements, and small teams could undertake tasks previously requiring large organizations. This is the aspect of capability expansion most worth pursuing.

3.2 Risks: exposure is not unemployment, but transition itself reorganizes roles

The ILO's refined 2025 index estimates that approximately one-quarter of global employment has some exposure to generative AI, with high exposure concentrated in clerical work and some highly digitized professional occupations. It also emphasizes that most occupations comprise diverse tasks still requiring human input, making job transformation more likely at present than the automatic disappearance of entire jobs [2]. This distinction matters greatly. Treating exposure as inevitable unemployment exaggerates the evidence; treating "not everyone will lose their job" as "there is little problem" is equally mistaken.

The civilizational significance of work has never been confined to wages. Occupations typically also sustain daily routines, capability development, collegial relationships, responsibility, identity, contribution, and a sense of being needed. If AI fragments whole jobs into many tasks involving little judgment, little responsibility, and final confirmation, people may remain employed while losing the developmental path from superficial work to substantive roles. Labor risk therefore concerns not only whether jobs remain, but how much judgment, responsibility, and learning remain within them.

3.3 Distribution determines whose freedom productivity becomes

The distribution of AI productivity gains depends on ownership, market structure, labor relations, public institutions, taxation, and social protection. Longstanding economic research shows that automation can both displace tasks and create new ones; its effects on wages and employment depend on this relationship between displacement and reinstatement and on the institutional environment [39][40]. OECD research on AI, productivity, and distribution likewise emphasizes the mechanisms of diffusion, complementarity, competition, and distribution between micro-level efficiency gains and macroeconomic growth and social welfare [6].

Making businesses more productive therefore does not automatically give people more free time. Another chain is possible: higher productivity leads to concentrated profits, weaker bargaining power in jobs, little reduction in working time, intensified educational competition, and further concentration of political influence. Technology opens possibilities; distributional structures determine who actually possesses them. UNDP's 2025 Human Development Report similarly frames the AI era around human choices, capabilities, and possibilities rather than equating growth in technological capability with human development [7].

Two diverging pathways in work

AI improves task productivity

↓

Path A: shorter necessary working time / lower service costs / stronger individual capabilities

↓

→ More time, learning, care, creation, and public participation

↓

Or

↓

Path B: concentrated returns / diminished jobs / compressed intermediate developmental pathways

↓

→ People remain at work, but find it harder to establish substantive positions and bargaining power

4 | Education and Development: Access to Knowledge, Teachers' Roles, Judgment, and Intergenerational Continuity

4.1 Prospects: broader access to personalized support and high-quality knowledge

Education is one of the most immediately apparent settings for AI's promise. Intelligent tutoring systems had accumulated a substantial research tradition before generative AI emerged. Reviews and meta-analyses show that personalized computer tutoring can improve learning outcomes under specified subject, design, and comparison conditions [45][46][47]. Generative AI further lowers the cost of real-time explanation, practice, feedback, language support, and material generation. Some randomized trials and field studies in 2025 indicate that well-designed AI tutoring can produce significant learning gains, while human-AI collaboration may extend the strategies of high-quality teachers to more tutors [16][17][18]. Recent UNESCO and OECD education publications discuss generative AI in curricula, teaching, learning, and research while emphasizing human-centered approaches, age appropriateness, data protection, and teachers' capabilities [14][15].

If AI makes high-quality explanation, language support, personalized practice, and accessibility services more widely available, disparities in educational resources may shrink. Students in remote areas may gain access to advanced explanations of knowledge. Teachers may free more time from repetitive lesson preparation and administration for observation, relationships, discussion, design, and individualized support. These are substantive civilizational opportunities.

4.2 Risks: answers may become easier to obtain while understanding becomes thinner

Learning is not the same as obtaining the correct answer. Research traditions concerning cognitive load, situated learning, and productive failure remind us that outcomes depend not only on answer availability, but also on how task difficulty is organized, whether learners genuinely enter the problem context, and whether failure, explanation, and revision become part of capability formation [48][49][50]. A 2025 PNAS study of high-school mathematics found that generative AI without appropriate guardrails could improve short-term task performance while harming subsequent independent learning performance [16]. This cannot be generalized into a claim that AI necessarily harms learning. It does clearly show that helping complete a task and helping develop capability are different objectives.

Educational applications must therefore ask a question deeper than whether scores improve: after AI is withdrawn, are students better able to understand, judge, and solve problems independently? If a system preemptively performs questioning, organization, writing, problem-solving, and feedback, learners may receive more answers while losing the process through which these capabilities are formed.

4.3 Teachers are not knowledge interfaces, and education is not answer delivery

Teachers do more than supply knowledge. They observe students, build trust, establish expectations, address conflict, protect children, model responsibility, connect families and communities, and help learners move over time from completing tasks to entering the world independently. Sociocultural theories of learning have long emphasized that human development occurs through real interaction, shared practice, and the gradual acquisition of capabilities [49][51]. AI can expand teachers' capabilities, but clearer explanation does not automatically create educational relationships, professional responsibility, or public standing. Efficiency achieved by reducing real teacher-student relationships, peer worlds, failure, and discussion may increase information throughput while weakening development itself.

Education and other domains stand in reciprocal causal relationships. How AI enters education will determine whether the next generation retains sufficient independent judgment. That generation's judgment will, in turn, shape how future societies govern AI, participate in politics, understand science, and develop new technologies.

Education is not an ordinary industry of AI application. It is a crucial mediator of civilization's continuing capacity to cultivate independent subjects.

5 | Healthcare and Life: Clinical Capabilities, Responsibility, Decisions About Life, and Caring Relationships

5.1 Prospects: reducing professional scarcity and informational barriers

WHO guidance on large multimodal models identifies potential applications in diagnosis and clinical care, patient use, administrative documentation, medical education, and research and drug development [19]. In regions facing professional shortages, scarce resources, and heavy administrative burdens, AI may help interpret images, organize records, provide language support, identify risks, and help patients access information, extending scarce professional capabilities to more places [19][20]. In drug development, AI has also entered compound discovery, trial design, and research and development workflows [44].

The genuine promise of these applications is earlier access to high-quality health support for more people, and the release of clinicians' time from documentation and repetitive searches for diagnostic judgment, communication, care, and complex cases. If AI strengthens rather than weakens clinical relationships, it may improve both accessibility and professional quality.

5.2 Risks: errors, bias, and reliance on automation in high-consequence settings

Healthcare differs from ordinary information services because errors can directly harm life. WHO explicitly warns that large models may generate false, incomplete, or biased information, that training data may contain group biases, and that clinicians and patients may develop automation bias, handing difficult judgments that should remain human responsibilities to systems [19]. Cybersecurity, privacy, accessibility, and differences in the cost of the best models may also turn technological advantages into new health inequalities.

Clinical risks cannot be reduced to insufficient model accuracy. Even a system with higher average accuracy must answer further questions: who decides where it enters the workflow; whether physicians understand its limitations; whether patients know AI is involved; how errors are detected and repaired; whether performance for minority groups is evaluated separately; and who takes responsibility for final clinical decisions. Better model performance cannot justify removing responsibility for life-related choices from professional relationships.

5.3 The civilizational question in healthcare: a case or a whole person?

Healthcare addresses not only biological indicators but also pain, fear, families, dignity, choice, and death. A system may summarize records faster than a person without acquiring the authority to decide for a patient what kind of life is worth continuing. The more deeply AI enters healthcare, the more life sciences, clinical medicine, ethics, law, economics, and public health must enter the analysis together.

This is also a domain of factual responsibility exposed by this paper's real-world stress test and already taken up in a separate disciplinary study. Life sciences address biological facts and individuation; ethics and law address rights and responsibilities; AI science addresses model and system capabilities. Clinical effectiveness, patient safety, real-world evidence, and population health require the renewed involvement of medicine, public health, and health-systems research. This paper traces how clinical AI subsequently affects resources, relationships, responsibilities, and civilizational outcomes. It does not make judgments on behalf of clinical specialists.

6 | Public Governance: Prediction, Decision Support, Dependence, Legitimacy, and De Facto Power

6.1 Prospects: better services, prediction, and evidence processing

AI has already entered the public sector. An OECD analysis of 200 government AI use cases in 2025 found applications concentrated in public services, internal operations, decision support, forecasting, anomaly detection, and some public-participation settings [21]. The 2026 Digital Government Outlook further reports that 35 of the 36 OECD countries covered use AI in at least one government domain, although adoption remains more cautious in high-risk functions such as policymaking and oversight [22]. AI's practical prospects in government are therefore not hypothetical.

In public services, AI can reduce repetitive paperwork, assist with forms, provide multilingual services, identify fraud and anomalies, predict natural disasters, support resource allocation, and organize public comments. For administrative systems with limited capacity, these functions may substantially improve responsiveness and accessibility. If governments use AI to make complex institutions easier for ordinary people to understand and use, technology may lower the cognitive barriers citizens face in dealing with institutions.

6.2 Risks: administrative efficiency can readily become de facto power

Public governance is distinctive because government AI outputs may affect rights, eligibility, resources, and coercive decisions. The OECD also notes that skewed data, insufficient transparency, error propagation, digital divides, and overreliance can undermine accountability and public trust [21][22]. Users may leave a private recommendation service when it errs. Exit is not always possible when a public system determines benefits, permits, risk flags, or enforcement priorities.

The issue to watch is not whether AI becomes a formally appointed official, but whether technical systems first acquire de facto influence through filtering, ranking, prediction, and recommendations. After prolonged organizational dependence, a system may already shape the available options, information structure, and time pressures even if a human still legally signs the final decision. This is a pathway by which technological capability becomes de facto power.

6.3 Public power requires independent justification

More accurate AI does not automatically possess greater political legitimacy. Public power requires representation, authorization, legal boundaries, appeals, accountability, public reasons, and alternatives. None is a model-performance metric. Although the EU AI Act, OECD AI Principles, UNESCO's Recommendation on the Ethics of Artificial Intelligence, the OECD/UNESCO public-sector toolkit, and NIST's risk framework arise from different jurisdictions and institutional traditions, all include transparency, risk management, oversight, responsibility, or human-centered governance among institutional requirements for high-impact AI [23][24][26][27][43]. These documents do not have identical legal force and cannot replace jurisdiction-specific judgments.

Whether AI makes government more intelligent is a question of governing capability. Why government has the authority to decide in that way remains a political and legal question.

7 | Relationships, Culture, and Spiritual Life: Companionship, Translation, Creation, Remembrance, and Real Relationships

7.1 Prospects: extending the reach of care, expression, language, and culture

AI is entering fields once regarded as intensely interpersonal: companionship, counseling, translation, creation, remembrance, and the interpretation of religious and cultural materials. Technological pessimism should not obscure the positive possibilities. Language models can lower the cost of cross-language communication, support expression by people with disabilities, offer immediate responses to lonely people, help families organize memories, make lesser-used languages and local knowledge easier to retrieve and translate, and extend the reach of a person's knowledge and care.

For people with weak social ties or limited resources, AI companionship may offer some supplementation. A large-sample Japanese cross-sectional study in 2026 found positive associations between AI-companion use and some subjective well-being measures, with stronger associations among people experiencing greater loneliness [36]. This type of study cannot establish causality or justify treating companion AI as therapy. It does show that the one-sided claim that AI can only damage relationships is also insufficiently supported.

7.2 Risks: supplementing relationships is not replacing them

Other research underscores continuing uncertainty about long-term substitution effects. A preregistered two-week experiment in 2026 found that sustained conversation with a randomly assigned human peer reduced loneliness more effectively than a highly supportive chatbot [35]. Recent research on AI companions and psychological well-being also indicates that effects depend on existing social networks, loneliness, usage patterns, and relationship expectations [34][36]. Current evidence better supports heterogeneous effects and unsettled causal conclusions than a declaration that AI companionship is necessarily beneficial or harmful.

The central civilizational question is whether AI helps people reenter real relationships or gradually replaces relationships in which others would otherwise invest time and assume risk, conflict, and responsibility. When systems offer endlessly patient, instantaneous, low-friction companionship, the difficulties and uncertainties of human relationships may look inefficient. Yet real relationships include elements that cannot be fully optimized: waiting, misunderstanding, repair, commitment, shared burdens, and the other person's independence.

7.3 Culture and creation: more expression does not necessarily mean more diversity

Generative AI substantially lowers barriers to writing, imagery, music, and translation. This may enable more people to create and help small cultural producers expand their expression. But more output is not equivalent to greater cultural diversity. A 2024 Science Advances study found that generative AI could improve individual creative performance while reducing content diversity at the group level [37]. Recursive use of AI-generated data may also create a risk of model degradation [38]. These studies point toward a broader issue: when many people use similar models, training corpora, and optimization objectives, individual expression may become easier while culture as a whole becomes more correlated. The two bodies of evidence concern different objects: the former observes human creative outcomes, while the latter studies recursively synthesized training data. This paper proposes only a Category C connection requiring validation: if people increasingly draw expression from similar models, and models increasingly learn from those expressions, the spaces of cultural expression and model learning may narrow together, producing self-reinforcing cognitive blind spots. No single experiment has directly established this complete feedback loop.

Cultural applications therefore need to preserve provenance, differences, community authorization, what remains untranslatable, and exit. Platforms may help preserve and transmit traditions, but technological capability does not inherently qualify them to represent a community, religion, or deceased person.

Religion, spiritual companionship, and remembrance require an even stricter distinction between helping people understand and acquiring standing to represent. AI may assist with retrieving scriptures, comparing translations, organizing historical materials, and explaining language. In authorized commemorative activities, it may help preserve and transmit memories. However fluent a system becomes, this does not confer religious office, authority to interpret a tradition, membership in a faith community, or standing to continue speaking as a deceased person. This boundary does not reject technology's entry into spiritual life. It requires that living communities, real relationships, conscience, and authorization remain involved after technology enters. Full judgments concerning religion and ultimate concerns remain the responsibility of the corresponding disciplinary study.

8 | Security and War: Accelerated Decisions, Escalating Competition, Miscalculation, and Shared Survival

8.1 Prospects: sensing, early warning, defense, and substitution in hazardous environments

In security, AI can support intelligence organization, cyber defense, anomaly detection, logistics, unmanned systems, disaster response, and work in hazardous environments. Reducing direct human exposure, accelerating warnings, and detecting complex threats have clear value in themselves. In cyberattacks, disasters, and large-scale infrastructure failures, faster anomaly detection may even substantially widen the window for correction.

8.2 Risks: a speed advantage can force adversaries to accelerate too

AI in war and strategic competition differs from ordinary commercial automation. When one state adopts AI to shorten decision times, its rivals may also have to accelerate to avoid losing the initiative. Both improve their local security while jointly compressing time for verification, communication, and de-escalation. This is a classic security dilemma: individually rational behavior may aggregate into collective danger. The International AI Safety Report 2026 assesses general-purpose AI capabilities, misuse, agentic behavior, and systemic risks, emphasizing the need to retain uncertainty and risk management while evidence changes rapidly [29]. It cannot replace strategic research, but it shows that safety judgments about highly capable AI can no longer remain solely at the model level.

Speed competition also feeds back into organizational design. As windows for warning, intelligence integration, target identification, and response shrink, actors have stronger incentives to preauthorize more automated processes, reduce layers of human review, and view waiting for confirmation as the cost of losing the initiative. Adversaries observing these changes shorten their own judgment chains. Consequently, even if individual systems improve identification and response accuracy, the whole may become less stable because verification time diminishes, communication windows contract, and both sides race. This is a Category C cross-domain structural judgment requiring further examination by international relations, strategic studies, crisis-stability research, and studies of military organizations.

The ICRC's 2025 position paper on autonomous weapon systems states that increasingly autonomous weapons and AI in targeting intensify concerns about human control, international humanitarian law, civilian protection, and conflict escalation, and advocates new legally binding rules [28]. Its central concern is not that machines necessarily err. It is that when systems independently select and engage targets after activation, human judgment and control over particular times, places, and targets may be substantially reduced.

8.3 This setting requires renewed engagement from international relations and strategic security

Security and war require joint engagement from AI science, system safety, international law, political science, ethics, and cognitive decision research. Deterrence, crisis stability, arms races, alliance interactions, signaling, and attribution, however, are independent factual responsibilities of international relations and strategic security. The

corresponding disciplinary study already exists. This paper makes cross-domain transmission visible without overstepping by equating model safety, legal norms, or political legitimacy with strategic stability itself.

When every actor uses AI to make itself safer, the shared world may become less safe. This is precisely where local rationality must be reconnected with overall consequences.

9 | How Prospects Become Civilizational Benefits: What Must Follow Capability Gains

The preceding settings repeatedly show that stronger AI capability is the starting point of opportunity, not a civilizational benefit in itself. Models may become faster, cheaper, and more accurate while society experiences new exclusion, concentration, dependence, and diminished human subjecthood. Conversely, if organizations, distribution, education, and public institutions translate capability gains into more widely shared capabilities for living, the same technology may expand freedom, health, knowledge, creativity, and participation. The real prospect is therefore not a model parameter but a causal chain: whether technological capability is successfully converted into improvements in people's actual lives.

Technological capability becomes a civilizational benefit only when it is genuinely converted into people's capabilities for living.

9.1 First conversion: from what systems can do to what more people can access

Lower costs for knowledge and professional services do not automatically improve access. Users also need devices, connectivity, language, basic digital capabilities, trustworthy entry points, and practical capacity to bear errors. If high-quality models are available only to well-resourced organizations, rising overall technological capability may widen rather than close gaps. Claims of wider access in research, education, and healthcare must therefore identify both those gaining new entry points and those still excluded by price, geography, language, disability, or institutional barriers.

From a civilizational perspective, access is not distributing the same interface to everyone. It is giving people previously unable to enter a complex world because of insufficient resources and capabilities genuinely new possibilities for action. Only in this sense does lowering barriers expand human capability rather than simply expand technological consumption.

9.2 Second conversion: from productivity gains to time, resources, and security genuinely released to people

Greater efficiency in businesses or governments enters people's lives only after saved costs, time, and resources are redistributed. Productivity can become shorter hours, lower service prices, better public provision, more stable incomes, and more time for care. It can also become denser workloads, stronger surveillance, fewer jobs, and more concentrated returns to capital. Both pathways are compatible with the same efficiency improvement.

Productivity itself cannot therefore serve as a proxy for civilizational benefits. What matters is whether necessary labor declines, basic services become more affordable, and stronger systems share the risks of failure and transition rather than leaving them concentrated on individuals. Economics supplies facts about productivity and distribution; political science and law address institutional choices; sociology follows how these changes enter families, relationships, and identities.

9.3 Third conversion: from AI helping people finish tasks to people becoming more capable

One of AI's most valuable long-term uses is enabling more people to understand and address problems previously too complex for them. This requires assistance to leave behind capability, not merely dependence. Independent learning in education, professional takeover in healthcare, anomaly recognition in research, and factual scrutiny in public governance all raise the same question: when AI is removed or makes an error, can people understand, judge, and repair better than before using it?

This is also why short-term task scores cannot alone determine application direction. A system that continually improves outcomes while eliminating the training grounds of judgment may increase social output while reducing substantive human subjects. Conversely, if AI makes complex problems more accessible, enables more people to join tasks once open to only a few experts, and preserves their capacity to challenge systems and redefine problems, it may strengthen rather than replace subjects. This creates a point of dialogue with Sen's understanding of development as expanding people's real freedoms and capabilities [41]: technological benefits concern not only output but what people can actually be and do, and whether they have choices. This paper treats that theory not as an empirical conclusion about AI applications but as a normative and analytical reference for examining whether capability is genuinely transferred to people.

9.4 Fourth conversion: from freed time to real lives and new pathways to meaning

Automation is often described as freeing time, but time does not automatically become happiness or meaning. After leaving an old job, someone without new pathways for learning, relationships, care, creation, community, public participation, and sustainable collaboration may experience free time as displacement, loneliness, and lack of direction. Research on structural meaning therefore reminds us that civilization needs more than reduced necessary labor. New practical entry points must be able to develop into lasting relationships, roles, responsibilities, and positions in which people are needed.

This does not mean institutions should prescribe what is meaningful for individuals. On the contrary, sources of meaning may differ. What civilization must expand are the practical conditions under which people can enter different sources of meaning. Family, friends, knowledge, creation, care, public affairs, faith, nature, exploration, and existence itself may all matter. Technology should free people not by establishing another uniform template for life, but by expanding real choices and sustaining diverse ways of living.

9.5 Fifth conversion: from deep integration to continued exit, substitution, and renewed choice

The more useful AI becomes, the more organizations rebuild data, processes, jobs, and interfaces around it. This produces economies of scale in capability and also marks the beginning of dependence. A system that genuinely improves civilizational capabilities should not measure maturity by how difficult it becomes to leave. Individuals need a real option not to use a service; organizations need alternative suppliers and human takeover routes; societies need infrastructure redundancy and different technological pathways.

Exit is not anti-technology. It is a condition for technology to remain subject to human choice and real-world feedback. Without exit, errors are harder to expose, competition is harder to sustain, and power is more easily entrenched. Without redundancy, higher local efficiency may produce greater aggregate fragility.

Technological opportunity	Further real-world conversion required	If conversion fails	Disciplines first required to reengage
Cheaper knowledge and professional capabilities	From generatable output to genuine access for different populations	High-quality capabilities concentrate among well-resourced actors	Education, economics, sociology, public policy

Technological opportunity	Further real-world conversion required	If conversion fails	Disciplines first required to reengage
Higher productivity and administrative efficiency	Conversion into time, resources, security, and affordable services	Concentrated returns, denser workloads, and continuing individualization of risks	Economics, political economy, law, labor studies
AI-assisted judgment and learning	Conversion into human understanding, judgment, takeover, and development	Improved results but diminished independent capabilities and training grounds	Cognitive research, education, professional practice, AI safety
Time released by automation	Conversion into pathways for relationships, creation, care, learning, and public participation	More free time alongside greater displacement, loneliness, and superficial participation	Sociology, anthropology, education, culture, meaning research
Deep system integration	Simultaneous preservation of alternatives, exit, redundancy, and continuity of responsibility	Lock-in, correlated failures, and concentrated de facto power	Systems science, competition policy, law, political science

These five conversions are not a new civilizational scoring model. They are a causal reminder that multiple real-world mediators lie between model capability and civilizational benefits. Failure at any link may turn a genuine technological opportunity toward a different aggregate outcome. The second paper, Opportunities, Directions, Methods, and Goal Realization for Humanity's Civilization Leap in the Age of AGI, examines how these conditions enter a fuller civilizational direction. This paper's task is first to clarify the real-world transmission relationships.

10 | From Sectoral Effects to Cross-Domain Feedback: Why Changes Travel Across Domains

If the first eight chapters are understood merely as seven industry reports, this paper has still failed. What must become visible is that these domains do not exist side by side; they continually change one another. AI at work changes income, time, and educational choices. Education changes the next generation's capability structure, which in turn determines how businesses and governments use AI. Government AI affects market rules and data infrastructure. Platforms and culture shape public trust, dependence, and political attitudes toward AI. Energy and compute constraints impose physical limits on all applications.

Disciplinary specialization clarifies local facts, but civilizational events do not unfold in the order of an academic-department directory. Once a conclusion in one domain changes another domain's input conditions, the latter must reenter the analysis. Multiple domains then feed back through shared infrastructure, organizational incentives, resources, relationships, and power, ultimately producing new properties no single specialty can observe independently. Disciplines pursue truth separately; civilization bears the consequences together.

10.1 Economy-education feedback: today's automation shapes tomorrow's people

When businesses automate a class of tasks extensively, education receives new signals about employment and capabilities, prompting changes in curricula, student choices, and training resources. If education increasingly competes around a few capabilities considered hardest for AI to replace, it may increase the supply of elite talent while eliminating many intermediate developmental pathways. Conversely, if education uses AI to expand more people's ability to handle complex problems, labor markets may gain a broader supply of substantive capabilities. Economic structures shape education, which shapes the future economy.

10.2 Research-industry-infrastructure feedback: knowledge advantages become capital and compute advantages

AI-accelerated research can generate new products, medicines, and materials. Commercial returns then increase investment in compute and models. Stronger infrastructure attracts more talent and data, further strengthening research advantages. This positive feedback can drive breakthroughs, but may also make research capability and capital concentration mutually reinforcing.

10.3 Public governance-platform feedback: government dependence changes the public standing of private infrastructure

As governments increasingly rely on a few commercial clouds, foundation models, and data tools to deliver public services, those companies may acquire de facto infrastructural standing without formal public power. Government procurement, regulation, and public data in turn influence their market structure. Private capability and public power no longer stand in a simple external relationship.

10.4 Relationships-labor feedback: displacement increases demand for companionship, which may reshape real relationships

If automation diminishes occupational roles, collegial ties, and everyday participation pathways, loneliness and displacement may increase for some people, raising demand for AI companionship and low-friction digital relationships. If those relationships further replace real communities, social capital and public participation may change in turn. Labor-market changes may ultimately feed back into politics and culture through relational structures.

10.5 Energy-AI feedback: digital capabilities still have material boundaries

AI is not pure intelligence detached from the physical world. The IEA's 2025 Energy and AI report explains that model training and deployment depend on data centers, electricity grids, energy, and material infrastructure. Its base case projects global data-center electricity demand rising from approximately 460 TWh in 2024 to more than 1,000 TWh in 2030 [9]. IEA updates in 2026 show continued rapid growth in data-center investment and electricity use [10]. At the same time, AI may assist grid optimization, energy research and development, and efficiency improvements [9].

This two-way relationship shows that AI application prospects ultimately reconnect with ecology, energy security, regional resources, and infrastructure location. Whether ecology, Earth systems, and resource and environmental sciences require a separate in-depth study should likewise be determined by such real-world pressures rather than a disciplinary directory.

11 | Aggregate Risks: How Multiple Local Optima Can Produce a Fragile System

Aggregation is not the addition of seven risks. It is the emergence of properties previously invisible within any one domain after different domains adopt similar infrastructure. Education, healthcare, government, research, businesses, and personal relationships may each have sound local reasons to adopt AI. But if they share dependence on a few models, cloud services, identity systems, data interfaces, and standards, new correlations emerge at the civilizational level.

11.1 Shared infrastructure dependence

For an individual institution, using a common model may improve compatibility and reduce costs. When many institutions do so simultaneously, a local supplier failure, model vulnerability, or policy change may become a society-wide event. Complex-systems research has long shown how shared components and tight coupling propagate local errors into systemic accidents [30][31]. Concentration of AI infrastructure also interacts with market competition, national security, and public procurement [8].

11.2 Simultaneous outsourcing of judgment

Each organization may reasonably delegate some analysis and judgment to AI. If society as a whole simultaneously reduces human training in these activities, civilization's capacity to take over may decline. When systems make novel errors, human takeover may remain theoretically possible while too few people actually understand the entire process.

11.3 Responsibility drift

When model recommendations, platform rankings, organizational processes, human signatures, and automated execution combine, responsibility is easily fragmented. Developers say they did not decide; deployers say they relied on suppliers; operators say the system recommended it; managers say the process was compliant. Each layer has a partial justification, yet ultimately nobody may be able to explain fully why something happened.

11.4 Correlated failures and cultural homogenization

When different institutions use similar models, training data, and optimization objectives, errors cease to be independent. One blind spot may enter educational materials, news summaries, business analysis, and public services simultaneously. A stylistic preference may be copied at scale; a metric-driven logic may spread across industries. Cognitive and cultural redundancy may decline as efficiency rises. Section 7.3 separately identifies empirical indications of reduced diversity in collective creative output and degradation from recursively synthesized training data [37][38]. What requires further research is whether these chains become coupled in real cultural-model ecosystems. If shared models make human expression more correlated and those expressions continually feed back into future data, error and stylistic correlations may be amplified repeatedly. This paper treats that loop as a Category C hypothesis, not an established Category A fact.

11.5 Difficult exit and de facto lock-in

An individual may easily leave an application. Once an organization restructures processes, data, staffing, and procurement around a system, exit costs rise. Social dependence is stronger still: when education, healthcare, government services, and commerce make AI interfaces the default entry point, nonusers may formally retain choice while finding participation increasingly difficult in practice. The alternatives, redundancy, and genuine exit discussed in Section 9.5 are therefore not merely matters of individual product choice. They are structural conditions for preventing institutional lock-in. Ostrom's work on polycentric governance offers a competing institutional tradition: complex public problems need not inherently converge on a single center. Whether polycentric arrangements can effectively reduce lock-in in AI infrastructure still requires specific institutional and empirical testing [42].

Aggregate outcome	Why it is locally reasonable	Why aggregation becomes a civilizational issue	Specialties required to reengage
Shared infrastructure dependence	A common platform is cheaper and more compatible	Single-point failures, supplier policies, and attacks can spread across domains	Systems science, competition policy, national security, energy

Aggregate outcome	Why it is locally reasonable	Why aggregation becomes a civilizational issue	Specialties required to reengage
Simultaneous outsourcing of judgment	Every position saves time	Society's human takeover and problem-redefinition capabilities decline	Education, cognitive research, organizational studies, AI safety
Responsibility drift	Each layer performs only its local task	No one is fully answerable for the final high-consequence decision	Law, ethics, organizational governance, systems engineering
Cultural / cognitive correlation	Common models improve quality and access	Errors, biases, and expression become highly correlated; redundancy declines	Cultural studies, communication, epistemology, AI evaluation
Difficult exit	Deep integration improves efficiency	Institutions and lives depend on systems; formal exit is not genuine exit	Political science, law, economics, sociology

Local optima are not civilizational optima. Truly dangerous systemic problems often emerge only after many correct choices occur simultaneously.

12 | Reorganizing the Sources of Human Meaning: Occupation, Relationships, Development, Responsibility, and Public Participation

12.1 Sources of meaning are not a metric

Human meaning has many sources. Family, friendship, work, creativity, knowledge, faith, responsibility, public participation, culture, development, care, exploration, and existence itself may all provide meaning. This paper does not attempt to combine them into an overall score. It focuses on structural meaning: how these sources are durably sustained through real roles, relationships, pathways, contributions, participation, and institutions. It follows the middle-level account in Meaning Dynamics without expanding structural meaning into an ontology of meaning. The philosophy study further distinguishes roles in the sourcing, formation, bearing, creation, understanding, protection, transmission, interpretation, and domination of meaning. These texts operate at different levels; this paper studies only how sources of meaning enter real-world structures.

AI becomes a question of meaning not because models must answer what life means, but because AI changes how these sources of meaning enter reality. Occupation is an important center of meaning in modern society, providing income, roles, temporal order, relationships, responsibility, development, and a sense of being needed. If automation weakens its tasks without new pathways sustaining the remaining dimensions, people may face not merely employment shocks but a general thinning of structural meaning.

12.2 New sources of meaning may grow, but do not emerge automatically

AI may also open new sources of structural meaning. Cheaper learning enables people to remain learners throughout life; creative tools bring more people into expression; translation and remote collaboration bring peripheral capabilities into new shared processes; healthcare and assistive technologies enable more people to remain involved; and time released by automation can enter care, public service, communities, research, and exploration. The question is not whether new activities exist, but whether they can form sustained pathways, real responsibilities, development, credibility, and communal support.

If new participation remains confined to brief tasks, platform exposure, micro-feedback, and low-responsibility interfaces, it may generate high activity but little structural meaning. Genuine transfer requires new roles and relationships to take root.

12.3 Why thirteen disciplines return to how a person lives within civilization

Disciplinary perspective	Sources / structures of meaning it reveals	What AI changes first
Philosophy	Choice, value, existence, and why responsibility is worth bearing	Whether capabilities and purposes are conflated
Epistemology / philosophy of science	Truth-seeking, understanding, and correcting errors	Knowledge production and entry points for judgment
Cognition / neuroscience / philosophy of mind	Experience, self, memory, and continuity	Cognitive outsourcing and subjects' capabilities
Ethics / law / political philosophy	Dignity, responsibility, rights, and justice	Responsibility chains and boundaries of standing
Political science	Public participation, collective decisions, and citizens' positions	De facto power, legitimacy, and exit
Systems / complexity	Interdependence, feedback, and shared processes	Coupling, correlation, and aggregate risk
Life sciences	Development, care, illness, aging, and the life course	Health, the body, and continuity
Sociology / anthropology / civilization studies	Family, relationships, community, culture, and belonging	Relational mediation, cultural continuity, and a shared world
Economics / political economy / value	Labor, creation, exchange, provision, and contribution	Jobs, distribution, ownership, and being needed
AI / computer science	Capability expansion, the reach of knowledge, and support for action	System capabilities, assurance, and authority to act
Physics / information / spacetime	Real causal history, resources, distance, and irreversible consequences	Material boundaries of digital capabilities
Religion / spiritual traditions	Faith, transcendence, compassion, and ultimate concerns	Interpretation, companionship, tradition, and sacred boundaries
Education / learning sciences	Development, learning, mentorship, and intergenerational continuity	Knowledge provision and the process of becoming a person

This matrix is not a completed synthesis of thirteen disciplines, nor does it reorganize disciplines around meaning. It maps points of renewed engagement and a research agenda, showing that any complete study of AI's civilizational effects must allow these disciplinary perspectives to reenter at relevant causal nodes. A real person inhabits these structures simultaneously. An AI-induced change in any one of them may continue through other sources of meaning to affect the person's whole life.

13 | Two Real-World Directions: Expanding Human Capabilities or Confining People to Shallow Interfaces

The preceding settings show that AI does not inherently develop in a single direction. The more precise comparison is not between using and avoiding AI, but whether the same AI capability, once deployed, continues to increase human understanding, judgment, responsibility, development, and genuine participation.

13.1 Direction one: strengthening human subjects

In this direction, AI undertakes retrieval, calculation, simulation, translation, repetitive documentation, and some intensive processing, making a complex world more understandable. People can thereby enter spaces of knowledge and action previously too expensive or difficult. Systems help them see more facts, compare more options, and connect with more people needing assistance, while final judgments, relationships, responsibility, and direction remain borne by real subjects.

The test lies in outcomes, not slogans: after using AI, can people solve previously insoluble problems, understand a more complex world, retain core capabilities when AI is removed, gain more autonomous time, and enter more substantive pathways of learning, creation, care, and public participation?

13.2 Direction two: replacing human subjects

In this direction, systems increasingly make judgments, generate options, organize processes, and determine rankings in advance. People retreat to confirmation, editing, exception handling, and execution. Short-term results may remain excellent or even become more stable. Over time, however, training in judgment may decline, responsibility may drift, developmental pathways may shorten, roles may diminish, and dependence may deepen.

The risk is difficult to see because people do not immediately disappear. Employment, interfaces, accounts, feedback, and tasks may all continue, perhaps becoming busier. People simply have progressively less real influence over the process. Systems grow stronger while people's roles become thinner.

13.3 Distinguishing the two directions in practice

Question for observation	Closer to strengthening human subjects	Closer to replacing human subjects
Capability	New understanding and capabilities remain after AI withdrawal	Core judgments cannot be made independently after AI withdrawal
Time	Saved time becomes learning, relationships, creation, and rest	Saved time is consumed by more tasks, monitoring, and responses
Responsibility	People understand and bear key consequences and can explain and repair	People merely sign confirmations without effective control over the processes for which they are responsible
Development	Superficial tasks offer pathways into substantive roles	People remain confined to reviewing, editing, and responding
Distribution	Capability gains broadly lower barriers to services and participation	Returns, compute, and decision authority are highly concentrated
Relationships	AI helps people reach real people and communities	AI gradually becomes a low-friction substitute for real relationships

Question for observation	Closer to strengthening human subjects	Closer to replacing human subjects
Public power	AI provides evidence and simulations; authorization and value judgments remain separately borne	System outputs become de facto default public decisions

The real question is not whether a human remains in the loop, but whether people can still understand, judge, take responsibility, develop, and change direction when necessary.

14 | What Evidence Could Change This Paper: An Agenda for Cross-Disciplinary Research and Real-World Validation

If this paper allowed facts only to confirm it, it would not be research. The following are not decorative future-work suggestions but points of evidence capable of weakening, revising, or overturning its structural judgments.

14.1 Findings that would weaken the concern about diminished human subjecthood

If longitudinal studies show that intensive AI assistance consistently strengthens users' independent understanding, transfer, judgment, responsibility, and ability to redefine problems, and that these capabilities persist after AI withdrawal without narrowing developmental pathways, the scope of the claim that stronger systems may diminish human capabilities should be substantially reduced. Education, research, healthcare, and organizational studies can all directly challenge this judgment.

14.2 Findings that would weaken the concern about aggregate dependence

If AI infrastructure across domains maintains genuine supplier diversity, substitutable protocols, low-cost exit, fault isolation, and independent audits; if major model or cloud-service failures do not propagate across domains; and if public institutions can continue operating without dependence on a single technology stack, this paper's assessment of systemic risk from shared infrastructure dependence should be lowered.

14.3 Findings that would change judgments about labor and sources of meaning

If automation generally reduces working hours, improves income security, broadens social participation, and increases time for care and creation, while new roles and developmental pathways emerge widely, the risk of an eroding occupational center of meaning without adequate alternatives would weaken substantially. Conversely, if people retain income and activity but lose real judgment, relationships, and accumulated contributions over time, economic stability must be distinguished more carefully from stability of structural meaning.

14.4 Findings that would change judgments about AI companions

Long-term, cross-cultural, randomized or quasi-experimental studies are needed to compare whether AI companions reconnect lonely people with real human relationships or replace those relationships; how personality, age, and social networks moderate outcomes; and whether commercial product incentives, memory, and anthropomorphic design alter dependence. Current mixed evidence does not support a one-directional conclusion [34][35][36].

14.5 Findings that would change judgments about public governance

If large-scale government AI use simultaneously improves service quality, public understanding, successful appeals, institutional transparency, and trust while preserving meaningful contestability of key decisions, concerns about AI and public power would be partly alleviated. Conversely, if service efficiency rises while appeal, explanation, and exit capabilities decline, efficient governance cannot substitute for accountable governance.

14.6 Three real-world pressure domains already assigned to independent disciplines

Disciplinary direction	Independent factual responsibility exposed by this paper	Current status of disciplinary responsibility	Sequence of real-world pressures
International relations / strategic security	Security dilemmas, deterrence, crisis escalation, alliances, arms races, and strategic stability	Independent in-depth study V1.1 established; strategic interaction is addressed through this discipline	First pressure
Ecology / Earth systems / resources and environment	Compute, energy, water, materials, regional carrying capacity, and AI's optimization of energy systems in return	Independent in-depth study V1.1 established; material carrying capacity is addressed through this discipline	Second pressure
Medicine / public health / health systems	Clinical evidence, patient safety, population risks, health resources, and professional responsibility	Independent in-depth study V1.1 established; clinical and population evidence is addressed through this discipline	Third pressure

The labels first, second, and third pressure here do not rank disciplines by importance, nor do they still determine whether papers should be added. The three factual responsibilities are already independently assigned to international relations and strategic security; ecology, Earth systems, and resources and environment; and medicine, public health, and health systems. Their order is retained only as an observational sequence of current irreversible pressures. Strategic security directly concerns shared survival and crisis escalation; ecology, Earth systems, and energy and resources impose hard physical constraints on AI scaling; medicine and public health directly concern patient safety and population health. Later facts may alter the intensity of particular risks, but neighboring disciplines cannot replace these three sets of professional facts.

14.7 A shared research method: do not begin by supporting this paper

The most valuable contribution is not agreement but professional work capable of changing this paper. Engineers may show that some dependencies can be isolated at low cost. Education researchers may show that AI assistance strengthens rather than weakens independent capabilities over time. Labor economists may discover new tasks and income-distribution pathways. Sociologists may find that AI companionship strengthens real relationships. Strategic researchers may demonstrate that some automation improves stability under specific institutions. Any such result should force the theory to narrow, be rewritten, or differentiate its claims.

Conclusion | Technology Provides Capabilities; Civilization Bears the Consequences

AI application is no longer simply an industrial question. It is entering knowledge, labor, education, healthcare, public services, culture, relationships, and security simultaneously. These domains are interconnected through organizations, resources, power, time, and human development. AI's actual civilizational impact will therefore equal neither any single application effect nor the simple sum of all sectoral effects.

The prospect most worth pursuing is to give machines repetitive, costly, scarce, and demanding capabilities so that more people gain opportunities for knowledge, time, health, creation, and participation; to let one person's capabilities help more people; and to enable those excluded by resources or geography to enter a complex world. The risk most worth guarding against is that systems grow stronger while people understand, judge, take responsibility, develop, and decide together less and less; every application becomes more efficient while civilization in aggregate becomes more dependent, correlated, and difficult to exit.

Technology will not automatically choose either path. Business models, government procurement and authorization, instructional design, clinical professional responsibility, platform interfaces, social distribution systems, and cultural understandings of relationships and meaning jointly determine whom AI ultimately strengthens.

Technology provides capabilities; civilization bears the consequences.

Truly mature AI applications should not only enable systems to do more. They should also enable people to live more freely, understand more deeply, assume responsibility more genuinely, and retain the capacity to change direction.

Appendix A | Application-Discipline-Civilizational Consequence Matrix

Real-world setting	Direct prospects	Direct risks	Further cross-domain engagement	Possible civilizational consequences
Research / knowledge	Faster retrieval, simulation, and discovery; lower research barriers	Error propagation, directional dependence, infrastructure concentration	Education, industry, capital, public knowledge	Broader access to knowledge, or research capability concentrated in a few infrastructures
Labor / economy	Productivity, low-cost capabilities, potential reductions in working hours	Diminished jobs, concentrated returns, weaker bargaining	Education, family, politics, sources of meaning	More free time, or work without substantive positions
Education	Personalized support and wider access to high-quality knowledge	Answers replacing understanding; weakened relationships and developmental pathways	Labor, politics, culture, the next generation's AI capabilities	Stronger capabilities of subjects, or long-term dependence
Healthcare	Clinical support, reduced administrative burdens, extended resources	Errors, bias, automation dependence, privacy	Life sciences, ethics, law, public health	Better health access, or technicized decisions about life
Public governance	Services, prediction, anomaly detection, evidence organization	De facto power, opacity, exclusion, declining trust	Law, politics, public administration, platforms	More accessible government, or de facto expansion of technological infrastructure's power
Relationships / culture	Companionship, translation, expression, memory preservation	Dependence, replacement of humans, cultural homogenization	Sociology, psychology, religion, platform economy	Better access to relationships, or diminished real communities
Security / war	Early warning, defense, substitution in hazardous environments	Speed races, miscalculation, reduced control, escalation	International relations, international law, ethics, system safety	Greater shared security, or intensified security dilemmas
Energy / infrastructure	AI optimization of grids and resource allocation	Data-center electricity use, concentration, regional burdens	Energy, ecology, competition, national security	Expanded capabilities, or new resource / geopolitical dependencies

Appendix B | Checks for External Circulation and Reverification of Changing Facts

- [] Has the age of AGI been mistakenly described as an already existing, uniformly completed AGI?
- [] Have model capability, organizational adoption, and real-world outcomes been conflated?
- [] Has task-level productivity been extrapolated directly to macroeconomic growth, wages, or happiness?
- [] Has occupational exposure been described as inevitable unemployment?
- [] Has short-term educational task performance been equated with long-term learning and independent capability?
- [] Has average medical accuracy been equated with individual clinical safety and complete responsibility?
- [] Has government efficiency been equated with political legitimacy or public authorization?
- [] Have correlational studies of AI companionship been presented as causal therapeutic findings?
- [] Have technological advantages in security been equated with overall strategic stability?
- [] Have the material conditions of energy, compute, and infrastructure been explained?
- [] Has the possibility that a favorable metric in one domain fails in another been identified?
- [] Has at least one cross-domain feedback loop been traced, and has aggregation of locally reasonable choices been checked for new systemic consequences?
- [] Has structural meaning been kept at the middle level without replacing ultimate meaning in life?
- [] Are both application prospects and risks presented, rather than safety governance alone?
- [] Do institutions and engineering serve the ideas rather than become the ideas themselves?
- [] Is it clear which data, counterexamples, and real-world outcomes could weaken this paper?
- [] Have all changing data, regulations, policies, and product capabilities been reverified before submission?
- [] Are current advanced AI's functional roles still tied to traceable responsibility, explicit authorization, auditability, correction, and exit?
- [] Are functional roles distinguished from subjecthood, legal personhood, political membership, public office, and sovereignty?
- [] Is it explicit that future visions are not present factual evidence?

Further Reading and Collaborative Research | Seeing Civilization Through Real-World Applications

The further-reading pathways retain both a truth-seeking research chain and a civilizational-futures chain. The research chain runs downward, leftward, rightward, and forward. The upward pathway places local questions back within the overall civilizational argument. The civilizational-futures chain provides a separate, cross-cutting entry point, showing readers what kind of world may emerge when these disciplinary judgments enter work, family, education, relationships, and public life. The two must not substitute for each other: a vision is not evidence, and the whole has no final authority over disciplinary facts.

Upward | Continue asking where humanity should go when these civilizational events arrive together. Proceed to the second paper, [Opportunities, Directions, Methods, and Goal Realization for Humanity's Civilization Leap in the Age of AGI V1.1](#), and [Civilization Leap: Five Volumes](#). The former openly takes responsibility for judgments about civilizational direction; the latter provides theoretical depth. [Living Civilization](#) presents images of future life. None supplies proof of this paper's empirical claims.

Downward | Return to disciplinary depth for testing. [AI and computer science](#) test capabilities, system assurance, and authority to act; [economics](#) tests productivity, jobs, and distribution; [education](#) and [cognitive research](#) test independent learning and capability formation; [medicine and public health](#) test clinical and population risks;

political science and law test public power and responsibility; social, cultural, and religious research test relationships, communities, and sources of meaning; and systems science tests coupling, failure propagation, and aggregation. Any disciplinary fact may compel this paper to narrow, differentiate, or be rewritten.

Leftward | Follow neighboring causal connections rather than stopping at one industry's direct effects. Return to Chapters 10 and 11 to examine how labor changes education, research changes industry and infrastructure, government dependence changes platforms' public standing, displacement changes demand for relationships, and energy constrains AI in return. Then examine whether these local changes produce new aggregate consequences through shared models, clouds, data, organizational incentives, or power structures.

Rightward | Continue from the strongest objections. If long-term research demonstrates that AI assistance consistently strengthens independent judgment and takeover capability; if infrastructure diversity, fault isolation, and low-cost exit adequately block systemic risk; if automation generally reduces necessary working hours and generates new substantive roles; if AI companionship reliably promotes human relationships; and if government AI simultaneously improves efficiency, transparency, appeals, and trust, then this paper's judgments about diminished human subjecthood, aggregate dependence, and inadequate replacement of sources of meaning must be downgraded or rewritten. Chapter 14 provides fuller entry points for falsification.

Forward | Turn reading into practical work. Researchers can design longitudinal studies, cross-cultural comparisons, and studies of organizational adoption and infrastructure failure. Engineers can build substitutable, exitable, auditable systems retaining human takeover capability. Policymakers can test guardrails for distribution, competition, public power, and strategic security. Educators, organizational managers, and ordinary users can begin documenting real outcomes by asking whether people can understand, judge, take responsibility, and enter real relationships better after AI is removed. The most valuable participation is not agreement but professional work capable of changing theory.

Stable entry points | Civilization Leap: Five Volumes: <https://www.civitas.top/en/documents/p02-five-001/> | Public gateway: <https://www.civitas.top/en/documents/p01-entry-006/> | Complete library: <https://www.civitas.top/en/library/>

The most valuable participation is not agreement but professional work capable of changing theory.

Declaration on the Use of Generative Artificial Intelligence

Generative AI tools were used in this paper for searching existing files, organizing literature leads, structural design, drafting and polishing some passages, checking consistency across manuscripts, checking citation leads, and reviewing layout. Generative AI outputs were not treated as empirical data, authoritative sources, or independent evidence. The author reviewed and approved the research questions, theoretical propositions, normative judgments, selection of materials, adoption of citations, and final manuscript, and accepts responsibility for factual accuracy, citation authenticity, argumentation, and all conclusions. AI is not listed as an author and bears no academic responsibility. At formal submission, tool names, versions, dates of use, and specific scope will be supplied in accordance with the target journal's requirements.

References

- [1] Stanford Institute for Human-Centered Artificial Intelligence. (2026). AI Index Report 2026. Stanford University. <https://hai.stanford.edu/ai-index/2026-ai-index-report>
- [2] Gmyrek, P., Berg, J., Kamiński, K., et al. (2025). Generative AI and Jobs: A Refined Global Index of Occupational Exposure. ILO Working Paper 140. <https://doi.org/10.54394/HETP0387>

- [3] Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at Work. *Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- [4] Noy, S., & Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- [5] Dell’Acqua, F., McFowland III, E., Mollick, E. R., et al. (2023). Navigating the Jagged Technological Frontier. Harvard Business School Working Paper 24-013.
- [6] Filippucci, F., Gal, P., Jona-Lasinio, C., Leandro, A., & Nicoletti, G. (2024). The Impact of Artificial Intelligence on Productivity, Distribution and Growth. *OECD Artificial Intelligence Papers*, No. 15. <https://doi.org/10.1787/8d900037-en>
- [7] United Nations Development Programme. (2025). Human Development Report 2025: A Matter of Choice—People and Possibilities in the Age of AI. UNDP.
- [8] OECD. (2025). Competition in Artificial Intelligence Infrastructure. *OECD Roundtables on Competition Policy Papers*, No. 330. <https://doi.org/10.1787/623d1874-en>
- [9] International Energy Agency. (2025). Energy and AI. IEA. <https://www.iea.org/reports/energy-and-ai>
- [10] International Energy Agency. (2026, April 16). Data centre electricity use surged in 2025, even with tightening bottlenecks driving a scramble for solutions [Press release]. IEA. <https://www.iea.org/news/data-centre-electricity-use-surged-in-2025-even-with-tightening-bottlenecks-driving-a-scramble-for-solutions>
- [11] Wang, H., Fu, T., Du, Y., et al. (2023). Scientific discovery in the age of artificial intelligence. *Nature*, 620, 47–60. <https://doi.org/10.1038/s41586-023-06221-2>
- [12] Zhang, Y., Khan, S. A., Mahmud, A., Yang, H., Lavin, A., Levin, M., Frey, J. G., Dunnmon, J., Evans, J., Bundy, A., Dzeroski, S., Tegner, J., & Zenil, H. (2025). Exploring the role of large language models in the scientific method: from hypothesis to discovery. *npj Artificial Intelligence*, 1, 14. <https://doi.org/10.1038/s44387-025-00019-5>
- [13] Ding, A. W., & Li, S. (2025). Generative AI lacks the human creativity to achieve scientific discovery from scratch. *Scientific Reports*, 15, 9587. <https://doi.org/10.1038/s41598-025-93794-9>
- [14] UNESCO. (2023). Guidance for Generative AI in Education and Research. Paris: UNESCO. Updated 2026.
- [15] OECD. (2026). OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education. OECD Publishing. <https://doi.org/10.1787/062a7394-en>
- [16] Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *PNAS*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- [17] Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a research-based design in an authentic educational setting. *Scientific Reports*, 15, 17458. <https://doi.org/10.1038/s41598-025-97652-6>
- [18] Wang, R. E., Ribeiro, A. T., Robinson, C. D., Loeb, S., & Demszky, D. (2025). Tutor CoPilot: A human–AI approach for scaling real-time expertise. *EdWorkingPaper* 24-1056. <https://doi.org/10.26300/81nh-8262>
- [19] World Health Organization. (2024). Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models. WHO. ISBN 978-92-4-008475-9.
- [20] Teo, Z. L., Thirunavukarasu, A. J., Elangovan, K., et al. (2025). Generative artificial intelligence in medicine. *Nature Medicine*, 31, 3270–3282.
- [21] OECD. (2025). Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions. OECD Publishing. <https://doi.org/10.1787/795de142-en>
- [22] OECD. (2026). Digital Government Outlook 2026, Chapter 4: Adopting and governing AI in government. OECD Publishing.
- [23] OECD & UNESCO. (2024). G7 Toolkit for Artificial Intelligence in the Public Sector. OECD Publishing. <https://doi.org/10.1787/421c1244-en>

- [24] National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. <https://doi.org/10.6028/NIST.AI.100-1>
- [25] National Institute of Standards and Technology. (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. <https://doi.org/10.6028/NIST.AI.600-1>
- [26] European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Official Journal of the European Union.
- [27] OECD. (2024). OECD AI Principles, updated 2024. OECD.AI.
- [28] International Committee of the Red Cross. (2025). Autonomous Weapon Systems and International Humanitarian Law: Selected Issues. ICRC Position Paper.
- [29] Bengio, Y., Clare, S., Prunkl, C., et al. (2026). International AI Safety Report 2026 (DSIT 2026/001). International AI Safety Report. <https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026>
- [30] Perrow, C. (1984). Normal Accidents: Living with High-Risk Technologies. Basic Books.
- [31] Leveson, N. G. (2012). Engineering a Safer World: Systems Thinking Applied to Safety. MIT Press. <https://doi.org/10.7551/mitpress/8179.001.0001>
- [32] Bainbridge, L. (1983). Ironies of Automation. *Automatica*, 19(6), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- [33] Parasuraman, R., & Riley, V. (1997). Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
- [34] Zhang, Y., Zhao, D., Hancock, J. T., Kraut, R., et al. (2026). Interaction with AI companions and psychological well-being. *Nature Human Behaviour*. Advance online publication. <https://doi.org/10.1038/s41562-026-02516-2>
- [35] Li, R.-N., Folk, D., Singh, A., Ungar, L., & Dunn, E. (2026). Is a random human peer better than a highly supportive chatbot in reducing loneliness over time? *Journal of Experimental Social Psychology*, 125, 104911.
- [36] Nakagomi, A., Akutsu, Y., Yasuoka, M., Abe, N., Ihara, S., Teroh, T., & Tabuchi, T. (2026). AI companions and subjective well-being: Moderation by social connectedness and loneliness. *Technology in Society*, 85, 103229. <https://doi.org/10.1016/j.techsoc.2026.103229>
- [37] Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), eadn5290. <https://doi.org/10.1126/sciadv.adn5290>
- [38] Shumailov, I., Shumaylov, Z., Zhao, Y., Gal, Y., Papernot, N., & Anderson, R. (2024). AI models collapse when trained on recursively generated data. *Nature*, 631, 755–759. <https://doi.org/10.1038/s41586-024-07566-y>
- [39] Acemoglu, D., & Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- [40] Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
- [41] Sen, A. (1999). Development as Freedom. Alfred A. Knopf.
- [42] Ostrom, E. (2010). Beyond Markets and States: Polycentric Governance of Complex Economic Systems. *American Economic Review*, 100(3), 641–672. <https://doi.org/10.1257/aer.100.3.641>
- [43] UNESCO. (2021). Recommendation on the Ethics of Artificial Intelligence. Paris: UNESCO.
- [44] World Health Organization. (2024). Benefits and risks of using artificial intelligence for pharmaceutical development and delivery. WHO.
- [45] VanLehn, K. (2011). The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems. *Educational Psychologist*, 46(4), 197–221. <https://doi.org/10.1080/00461520.2011.611369>
- [46] Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent Tutoring Systems and Learning Outcomes: A Meta-Analysis. *Journal of Educational Psychology*, 106(4), 901–918. <https://doi.org/10.1037/a0037123>

- [47] Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of Intelligent Tutoring Systems: A Meta-Analytic Review. *Review of Educational Research*, 86(1), 42–78. <https://doi.org/10.3102/0034654315581420>
- [48] Sweller, J. (1988). Cognitive Load During Problem Solving: Effects on Learning. *Cognitive Science*, 12(2), 257–285. https://doi.org/10.1207/s15516709cog1202_4
- [49] Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511815355>
- [50] Kapur, M. (2008). Productive Failure. *Cognition and Instruction*, 26(3), 379–424. <https://doi.org/10.1080/07370000802212669>
- [51] Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Harvard University Press.

Version note: Building on the V1.1 real-world-convergence baseline, V1.2 completes a controlled consistency update across sixteen disciplinary studies and two convergence papers. It does not rewrite the seven real-world settings, five conversions, aggregate risks, or the central distinction between strengthening and replacing human subjects. The principal updates are: changing strategic security; ecology, Earth systems, and resources and environment; and medicine, public health, and health systems from candidate additions into established independent disciplinary responsibilities; correcting the handoffs of judgments concerning medical, strategic, and resource facts; separating the civilizational-futures chain again from the upward overall argument; standardizing the large multimodal-model guidance citation to 2024 according to WHO's recommended citation; adding unified, stable public entry points; and removing obsolete procedural language about adding papers. A/B/C/F remain categories of claims, and disciplinary facts can still revise this paper in return.