

Opportunities, Directions, Methods, and Goal Realization for Humanity's Civilization Leap in the AGI Era

Why We Must Change, Where We Can Go, and How to Get There

AGI's greatest opportunity for humanity is not simply access to a more powerful tool. For the first time, it may become possible to reorganize production, knowledge, time, relationships, and the human place within them.

In-Depth Synthesis on Civilizational Direction V1.1

Author: Zijun Fu

Produced by: Civilization Leap Research Group

Evidence cutoff: August 26, 2026

Abstract

Artificial intelligence is expanding from individual tools into a general-purpose capability condition spanning scientific research, labor, education, medicine, governance, culture, relationships, security, and infrastructure. The 2026 AI Index reports that AI adoption among surveyed organizations has reached 88%, with 79% using generative AI in at least one business function. Meanwhile, real-world deployment of agentic systems still lags significantly behind demonstrations of model capabilities [1]. The International Labour Organization's 2025 task-level study estimates that roughly one-quarter of global employment falls within some degree of generative AI exposure, but at present the better-supported outcome is job transformation rather than the straightforward disappearance of entire jobs [3]. The IEA projects that, in its base case, global data center electricity consumption will rise from its 2024 level to approximately 945 TWh in 2030, with AI an important driver of growth [9]. These facts do not prove that any particular future is inevitable. They do, however, suffice to show that AI is no longer merely an industry-specific variable: it is becoming a civilizational condition confronting multiple disciplines, institutions, and real-world actors together.

Building on the analysis of real-world transmission mechanisms in *AI Application Prospects and Risks in the AGI Era: From Capability Expansion to Civilizational Consequences*, this paper does not repeat its seven application domains. Instead, it asks why humanity must actively choose a civilizational direction when these changes occur together; which changes are already visible, which are highly probable if current trends in capabilities, costs, and organizational adoption continue, and which futures must remain open; and what civilizational opportunities AGI and its precursor technologies actually open for humanity, alongside the undesirable directions toward which they might otherwise propel civilization.

This paper explicitly distinguishes four argumentative statuses: A, B, C, and F. Existing adoption, labor exposure, energy constraints, and formal international principles constitute category A facts or institutional materials. Competing explanations of long-term educational, organizational, and social mechanisms belong to category B. Cross-domain judgments about the future, such as “judgment becoming scarcer,” “the reorganization of sources of occupational meaning,” and “the possible formation of systemic dependence,” belong to category C. Claims such as “continued existence is the baseline, while freedom from survival insecurity → happiness in life → a more meaningful life defines civilization’s progressive direction,” and “technology should extend human capabilities rather than silently eliminate the human position as a subject,” are category F normative commitments for which the author openly takes responsibility. The paper neither disguises civilizational direction as scientific fact nor allows facts automatically to replace value choices.

The paper argues that the most important opportunity in the AGI era is not “having machines do more for people,” but translating new productive, knowledge-related, scientific, educational, medical, and collaborative capabilities into people’s actual capabilities for living. This means making basic subsistence progressively less dependent on compulsory attachment to a single job; enabling more people to have discretionary time, health, safety, relationships, and real choices; and widening access to diverse sources of meaning, including creation, care, learning, public participation, faith, exploration, and being needed. Conversely, if productivity growth fails to translate into widely shared freedom; if organizations continually outsource judgment and responsibility; if education becomes answer delivery; if infrastructure and public power become highly concentrated; and if relationships and culture are increasingly mediated by a small number of similar systems, civilization may make its systems stronger while diminishing the depth of human subjecthood.

Civilization Leap is therefore understood here as a directionally guided reorganization of reality, rather than a one-time replacement of institutions. Its method emphasizes starting with real problems and genuinely affected people; allowing old and new arrangements to coexist; comparative validation; and the possibility of failure and exit. AI should expand human understanding, judgment, creativity, responsibility, and the collective formation of

direction. Priorities should change across phases, without making emergency arrangements from the launch phase permanent. Goal realization does not involve constructing an “overall civilization score,” but observing whether three kinds of directional change continue to emerge. The paper’s ultimate proposition is this: technology opens possibilities, thought determines direction, and reality tests the path.

Keywords: AGI era; Civilization Leap; artificial intelligence; civilizational event; human development; structural meaning; capacities for subjecthood; freedom from survival insecurity; happiness in life; meaning in life

Abstract

Artificial intelligence is becoming a general-purpose condition that cuts across science, work, education, health care, government, culture, relationships, security, and infrastructure. This paper does not attempt to predict an exact arrival date for AGI, nor does it assume machine consciousness, superintelligence, or post-scarcity. Instead, it asks a more immediate civilizational question: as increasingly capable AI systems enter real institutions and reshape knowledge, labor, responsibility, resources, relationships, and power, what directions remain open to human choice?

The paper distinguishes what is already visible, what may plausibly continue if current capability, cost, and adoption trends persist, and what remains genuinely open. It treats the following as normative commitments rather than scientific findings: civilizational survival is a non-negotiable floor; beyond survival, the desired direction is progressively toward freedom from survival insecurity, a happier and richer lived life, and wider conditions for people to generate meaning through relationships, creation, care, learning, responsibility, public participation, faith, exploration, and existence itself. These goals cannot be collapsed into one score or traded against basic rights.

The proposed method is deliberately non-blueprint: begin with real problems and affected people; allow old and new arrangements to coexist; compare outcomes; preserve the capacity to fail, exit, and revise; use AI to augment human judgment, responsibility, growth, and direction-setting rather than silently displace them; and adapt priorities across phases. Civilizational direction is therefore not a forecast. It is a choice that must remain answerable to evidence, professional knowledge, lived experience, and failure.

Keywords: AGI era; Civilization Leap; artificial intelligence; human development; structural meaning; human agency; well-being; meaningful life

Research Positioning, Scope, and Argumentative Discipline

This is the second of two synthesis papers on moving from disciplinary inquiry to the civilizational level. The first, *AI Application Prospects and Risks in the AGI Era: From Capability Expansion to Civilizational Consequences V1.2*, traces AI applications from model capabilities through organizations, human-machine divisions of labor, resources, power, and relationships to civilizational consequences. The present paper draws only on the real-world conditions that the first paper establishes, rather than repeating its application evidence chapter by chapter. Its responsibility is to answer directly, once these real-world changes converge into civilizational events, where humanity wants them to lead, how technological opportunities can be translated into better conditions of life, and how we can tell whether we are actually getting there.

Reality unfolds through causation, not through disciplines. Disciplines pursue truth separately; civilization bears the consequences together.

These two statements do not call for all disciplines to merge into a “super-discipline.” Computer science remains responsible for capabilities and system assurance; economics for production, distribution, and incentives; the life

sciences and medicine for facts about life and evidence about health; educational research for learning and development; law and political science for standing, authorization, rights, and public power; and philosophy and ethics for the open discussion of purposes and norms. The civilizational scale simply continues tracing how these separately established judgments interact, generate feedback, and aggregate once they enter the same reality, ultimately changing people's lives.

Argumentative status	Typical content in this paper	What it can support	What does not follow automatically
A Established facts / formal materials	AI adoption, employment exposure, energy demand, formal international documents, existing institutional conditions	Factual judgments within specified times, objects, and conditions	Complete social consequences, ultimate direction, or subjecthood
B Frontier / competing explanations	Long-term educational effects, organizational incentives, capability retention, effects on relationships	Competing mechanisms and research directions	Presentation as already convergent universal conclusions
C Cross-domain structural judgments	Highly probable continuation scenarios, reorganization of meaning sources, capability concentration, and de facto power	Refutable civilizational hypotheses	Claims that a single discipline has already proven them
F Philosophical / normative commitments	The three goals, the human position as a subject, and the prohibition on eliminating basic rights for efficiency	An open account of the direction this paper chooses and why	Replacement of empirical research, legal procedures, or public authorization

This paper uses “AGI era” to describe advanced AI’s continuing entry into knowledge production, organizational action, resource allocation, and high-consequence decision chains, prompting society to begin reorganizing around the general availability of more powerful intelligence. It is not a factual presupposition that a fully realized and uniformly defined AGI entity already exists. Present systems are described with specific terms such as “current advanced AI,” “generative AI,” and “agentic systems”; future AGI is treated as an object of prospective inquiry.

Current public evidence remains insufficient to establish that advanced AI possesses subjective experience, non-transferable continuity of subjecthood, or awareness of consequences for itself. Capability, experience, continuity of subjecthood, protection, legal standing, political membership, public authority, and sovereignty each require separate justification; they cannot be linked automatically. If sufficiently strong evidence of subjecthood emerges in the future, converging across multiple methods, its ethical and institutional status should be reconsidered. This paper neither uses today’s insufficient evidence to close the future permanently nor uses future possibilities to authorize present systems in advance.

Contents

Introduction | Why a Genuine Window for Civilizational Choice Is Opening for the First Time

1 | Already Visible: Why AI Is Becoming a Civilizational Condition Facing All Disciplines

2 | Highly Probable Possibilities: What Is More Likely If Current Capability, Cost, and Adoption Trends Continue

3 | Genuinely Open Futures: Major Questions That Cannot Yet Be Declared Inevitable

4 | Opportunities: From More Powerful Tools to Reorganizing the Conditions of Human Life

5 | What Cannot Happen Automatically: Technological Progress Does Not Produce a Good Civilization by Itself

6 | Direction: The Baseline of Continued Existence, Freedom from Survival Insecurity, Happiness in Life, and a More Meaningful Life

7 | The Human Place: Civilization Leap Means Real Lives and Diverse Sources of Meaning, Not Merely “Something Left to Do”

8 | Shared Problems: How Disciplines Can Address Civilizational Events Together Without Erasing Professional Boundaries

9 | Method One: Start with Real Problems and Genuinely Affected People

10 | Method Two: Allow Old and New Arrangements to Coexist, Compare and Validate, and Permit Failure and Exit

11 | Method Three: Use AI to Expand Human Capabilities While Retaining Judgment, Responsibility, Development, and Direction

12 | Method Four: Change Priorities by Phase Without Making Emergency Arrangements Permanent

13 | Goal Realization: Observing Freedom from Survival Insecurity Becoming a Reality

14 | Goal Realization: Observing Happiness in Life Becoming a Reality

15 | Goal Realization: Observing the Expansion of Conditions for a More Meaningful Life

16 | From Direction to the Lifeworld: The Responsibilities of Vision, Institutions, Engineering, and Real-World Systems

17 | How Reality Can Refute Us: Failure, Counterexamples, and Civilizational Self-Correction

Conclusion | Technology Opens Possibilities, Thought Determines Direction, and Reality Tests the Path

Appendix A | Civilizational Event–Discipline–Direction Matrix

Appendix B | Already Visible, Highly Probable Continuation, and Genuinely Open: A Checklist for Judgments About the Future

Further Reading and Joint Participation | Letting Reality Continue to Change the Theory, Not Merely Finishing It

Generative AI Use Statement

References

Introduction | Why a Genuine Window for Civilizational Choice Is Opening for the First Time

Many technologies have changed civilization. Agriculture changed settlement, land, and power; industrialization rewrote labor, cities, education, and time; the internet changed knowledge, communication, and organization. AI's distinctive feature is not that "machines are helping people for the first time." It is that an increasing number of cognitive, generative, analytical, judgment-preparation, coordination, and execution tasks once requiring humans can now be accessed across domains through the same general-purpose technological infrastructure. The 2026 AI Index records continued rapid growth in both organizational adoption and individual use, alongside AI's widening presence in science, education, medicine, and government [1]. This change is still far from complete, but it has already altered the conditions under which many domains can address their problems independently.

Looking only at individual applications makes it easy to interpret the present as "another efficiency revolution." Yet when AI simultaneously enters work, knowledge, education, medicine, governance, relationships, culture, and interstate competition, the same technological condition crosses economic, legal, social, political, psychological, biological, and meaning-related structures. A firm's decision to automate may be locally rational; many firms automating simultaneously will change labor bargaining, educational choices, and public finances. A school's adoption of AI tutoring may improve learning efficiency; the long-term restructuring of an entire education system around AI will change the next generation's judgment and developmental pathways. A government's use of AI may accelerate services; deep dependence across multiple public systems on a few infrastructures will raise questions of de facto power, capacity to exit, and system redundancy.

This is what a "window for civilizational choice" means. It does not mean that humanity suddenly has a blank blueprint or that old institutions will automatically collapse. Rather, a set of conditions once difficult to change simultaneously—productive capacity, knowledge costs, organizational coordination, educational provision, scientific discovery, medical assistance, and cross-language communication—is being loosened by the same technology at the same time. The more powerful technology becomes, the less civilization can limit itself to asking "Can we do it?" It must ask anew: "What is it for? Who decides? How do the gains enter people's lives? Which processes must not be silently handed over?"

UNDP's Human Development Report 2025 describes the AI era as "a matter of choice," emphasizing that human development concerns not how humanlike machines appear, but the choices and capabilities people can actually possess and realize [2]. This paper shares that international development framework's concern. However, responsibility for the three goals and Civilization Leap's normative direction rests with the author; neither UNDP nor any other international institution is invoked to authorize them normatively.

Technology offers opportunities; civilization determines what those opportunities ultimately become.

The following discussion distinguishes three layers of judgments about the future. The first concerns realities already visible, without waiting for AGI. The second concerns highly probable changes with substantial real-world grounds, conditional on the continuation of current capability, cost, and organizational adoption trends, and still open to refutation. The third concerns genuinely open futures: when AGI might meet some unified definition, whether machines will develop subjecthood, whether superintelligence will emerge, and whether extreme abundance will be realized. This distinction prevents civilizational direction from being held captive by technological determinism, while ensuring that future uncertainty does not deprive us of reasons to act today.

1 | Already Visible: Why AI Is Becoming a Civilizational Condition Facing All Disciplines

1.1 A civilizational event is a shared change in multiple domains' input conditions, not merely "major news"

A "civilizational event" here is not a product launch or one company taking the lead. It is a change that simultaneously alters the input conditions of multiple disciplines and social systems, producing consequences through feedback that no single domain can bear independently. AI's entry into knowledge work is one such event: it first changes task costs and speeds, then changes job design, educational signals, infrastructure investment, organizational power, legal liability, and human developmental pathways.

Such events do not matter equally to every discipline, but no profession can treat them as "peripheral background." Reality does not wait for economics, political science, education, and philosophy to finish their research in sequence: these changes interact simultaneously. Professional boundaries still protect the quality of truth-seeking, but can no longer be treated as the endpoints of real-world causal chains.

1.2 Eight kinds of shared change already under way

Civilizational change	Already visible reality	Disciplines that must engage together
Intelligent capabilities becoming widely available across society	More individuals and organizations can directly access generation, analysis, coding, translation, and planning capabilities	AI / computing, economics, education, organizational studies, society and culture
Restructuring of work tasks and organizational processes	Task exposure and workflow automation expand; at present, changes in job content are more likely than uniform job disappearance [3]	Economics, organizational studies, education, law and politics, sociology
Changes in knowledge production	AI enters research, literature work, code, experimental design, and scientific infrastructure [1][19]	Epistemology, philosophy of science, natural sciences, education, political economy
Progressive externalization of judgment	Recommendations, decision support, automation, and agents route more stages of judgment through systems first	Cognitive science, systems science, ethics and law, politics, education
Expansion of computing and infrastructure	AI computing capacity, data centers, energy, and cloud infrastructure become new physical and capital conditions [9][10]	Energy / Earth systems, economics, national security, competition policy
AI's emerging use in public governance	Growing institutional adoption of AI in public services, policy analysis, and administrative processes [20][36]	Politics, law, public administration, epistemology, AI
Technological mediation of relationships, culture, and education	AI enters learning, expression, companionship, translation, and cultural production, changing how people encounter knowledge and one another	Education, sociology and anthropology, cognitive science, religion and culture, ethics
Interstate competition and shared risks rising together	AI capabilities become increasingly linked to economics, security, and technological sovereignty, while risks spread across borders [11][16]	International relations / strategic security, politics, economics, systems science, law

1.3 “Every discipline should see this” does not mean “every discipline should study the same thing”

Computer science is not obliged to measure income distribution for economics; economics cannot assess long-term learning for educational research; philosophy cannot establish treatment facts for medicine; and civilizational theory certainly cannot judge strategic stability for international relations. Confronting these changes together means that every discipline must know whose causal chain its conclusions will subsequently enter, and when changes elsewhere will in turn alter its own object of study.

Confronting a problem together does not mean adjudicating it together. What is integrated is the real-world problem and its causal chains, not the disciplines’ standards of evidence.

2 | Highly Probable Possibilities: What Is More Likely If Current Capability, Cost, and Adoption Trends Continue

“Highly probable” here denotes neither statistical confidence nor a predicted timetable. It is a category C conditional judgment: if current AI capabilities continue improving, usage costs continue falling, and organizational adoption continues expanding, without major technical, legal, economic, or social resistance sufficient to reverse these trends, there are strong real-world reasons for the following developments to continue. Each includes conditions that could change the judgment.

2.1 Knowledge costs continue falling while judgment grows in importance

The costs of generation, retrieval, translation, summarization, coding, and preliminary analysis have already fallen substantially. New AI infrastructure is also forming within scientific research [1][19]. If this trend continues, what becomes scarcer may no longer be merely “having answers,” but asking worthwhile questions, distinguishing stronger from weaker evidence, handling conflicting goals, recognizing anomalies, and taking responsibility for outcomes.

This is not an ontological declaration that “human judgment can never be replaced.” If future systems reliably redefine problems, evaluate evidence, handle value conflicts, and bear consequences in open environments, and society is willing to accord them the corresponding standing and responsibility, this judgment must be reopened. For now, it simply reminds education and organizations not to eliminate the processes through which judgment develops as answers become cheaper.

2.2 Many occupations are more likely to be restructured first than to disappear uniformly

The ILO’s 2025 global index estimates that roughly one-quarter of employment is exposed to generative AI to some degree. It explicitly states that transformation is more likely than wholesale redundancy because most occupations still include tasks requiring human input [3]. Experimental and field studies have also observed task-level productivity gains in some kinds of knowledge work [4][5].

The highly probable issue is therefore not simply “Will there still be jobs?” but what remains inside those jobs: whether people still encounter complete problems, accumulate experience, exercise judgment, assume responsibility, and move from shallow tasks into substantial roles. A job that persists while its developmental pathway disappears can still constitute displacement at the civilizational level.

2.3 Organizational dependence will change living environments faster than individual preferences

Once AI is embedded in the standard procedures of businesses, schools, hospitals, governments, and platforms, whether an individual actively chooses to use it is no longer the only variable. Recruitment, evaluation, education, services, credit, public administration, and information environments may all be changed by organizational AI adoption. The OECD–UNESCO public-sector toolkit, OECD research on government AI, and NIST’s risk frameworks all treat AI as a matter of the entire organizational lifecycle and governance, rather than merely a personal tool [17][18][20][36].

One of the most important future decisions will therefore occur at the organizational level: who receives the time saved, whether AI outputs become the default, who retains powers of veto and appeal, and whether human training and redundancy are preserved.

2.4 Productivity may improve, but the distribution of gains will not improve automatically

Existing experimental and field research supports productivity gains for certain tasks and settings [4][5]. OECD macroeconomic research also warns that the movement from micro-level efficiency to macro-level productivity, wages, and welfare still passes through diffusion, competition, complementarities, and distribution mechanisms [6]. Research on automation has long shown that technology can displace tasks and also create new ones [7][8].

There is therefore no automatic equivalence between “AI raises productivity” and “people gain more free time.” If ownership, market structure, and bargaining relations concentrate the gains, productivity growth can coexist with employment insecurity, educational competition, and concentrated power.

2.5 Education will be forced to reconsider its purpose

As explanations, practice, writing, code, and answers become more accessible, education organized primarily around transferring knowledge and producing standard answers will find it increasingly difficult to justify its value. UNESCO already calls on countries to center human needs, age appropriateness, data protection, and capability development in responding to generative AI in education [14]. Educational research also shows that AI tutoring can help learning, but that without guardrails, improved short-term performance can coexist with impaired independent learning [29][30].

The highly probable change is therefore not “teachers disappearing.” It is education having to answer anew which knowledge people must master themselves, which processes they must experience firsthand, and how AI assistance can support stronger understanding, judgment, responsibility, relationships, and self-development.

2.6 Technological infrastructure will increasingly become a foundation of de facto power

Frontier AI requires computing capacity, chips, data centers, energy, networks, and cloud services. The OECD has made competition in AI infrastructure a distinct competition-policy issue [10]. In its base case, the IEA projects data center electricity demand at approximately 945 TWh by 2030, with AI among the main growth drivers [9]. The more technological capability underpins public and economic activities, the more likely those controlling its infrastructure are to gain real-world influence beyond that of traditional product suppliers.

De facto power is not legitimate power. The importance of infrastructure cannot automatically generate political authorization, public authority, or sovereignty. It will, however, compel political science, law, competition policy, international relations, and energy-systems research to engage together.

2.7 Sources of human meaning will be reorganized from a single occupational center toward multiple sources

In modern society, occupations have long bundled income, identity, daily schedules, colleague relationships, development, contribution, and being needed. If AI keeps restructuring jobs, this “bundled structure of meaning” will loosen. The risk is that the old center weakens before new supports form; the opportunity is that creation, learning, care, family, community, public participation, culture, faith, and exploration gain greater practical scope.

This “reorganization of sources of meaning” is a category C structural judgment, not a prediction of everyone’s psychological state. If long-term research shows that occupational structures remain stable supports for identity, development, and being needed even under extensive automation, or that new non-occupational pathways struggle to sustain lasting roles, this judgment should be narrowed.

2.8 Interstate competition will accelerate while shared risks also increase

When AI has clear economic, scientific, and security value, states and firms alike have reasons to seek leadership. The 2026 AI Index already shows the growing importance of interstate model-capability competition, AI sovereignty, and infrastructure competition [1]. The United Nations’ Global Digital Compact explicitly locates AI governance within international cooperation, human rights, development, capacity-building, and global risk [11].

Security presents a classic pattern of local rationality producing collective risk: one side accelerates to avoid falling behind, forcing others to follow. As decision times shorten, incentives for automation and pre-authorization intensify, potentially compressing opportunities for communication and correction further. The International AI Safety Report 2026 also notes the coexistence of misuse, malfunction, and systemic risks in general-purpose AI, alongside significant limitations in existing safeguards [16]. The ICRC’s position on autonomous weapons further emphasizes the need for explicit boundaries concerning human control, legal responsibility, and unacceptable risk [35]. Full judgments of strategic stability must still be undertaken independently by specialists in international relations and strategic security.

3 | Genuinely Open Futures: Major Questions That Cannot Yet Be Declared Inevitable

Civilizational direction requires foresight, but cannot gain strength by presenting unknowns as known. The following questions could profoundly change civilization, yet remain unsuitable as inevitable premises for all of today’s institutional choices.

Genuinely open question	Why it cannot be settled in advance	What can still be done today
When will AGI emerge in a unified sense?	Capability definitions, evaluations, and system boundaries are still changing; no single date can be treated as fact	Govern capabilities and deployments already present, while retaining prospective stress tests
Will AI develop subjective experience or continuity of subjecthood?	Current public evidence is insufficient, and the relevant scientific theories themselves have not converged	Remain open to evidence without letting functional roles automatically confer standing as a subject
Will sustainable superintelligence or uncontrollable systems emerge?	Capabilities are advancing rapidly, but open-world reliability, controllability, and safety assurance remain substantially uncertain [16]	Develop layered safety, independent evaluation, social resilience, and designs that permit stopping

Genuinely open question	Why it cannot be settled in advance	What can still be done today
Will extreme material abundance be achieved?	Productivity, energy, resources, ownership, distribution, and ecological constraints may all become bottlenecks	First convert existing productivity gains into more widely shared capabilities for living
Mind uploading, digital immortality, and post-biological civilization	These involve unresolved questions of consciousness, identity, life, and physics	Keep them open as long-term research, without substituting them for the rights and realities of present finite lives

This distinction has another important consequence: Civilization Leap need not wait for every open question to be resolved. Once it is established that AI capabilities are increasing, organizations are becoming more dependent, and the conditions of production and knowledge are changing, humanity must decide how these changes enter distribution, education, public power, and daily life. Waiting for “real AGI” before discussing civilizational direction would instead allow organizational inertia to settle many directional choices in advance.

Many technological trends may be difficult to reverse, but choices remain about how they enter civilization.

4 | Opportunities: From More Powerful Tools to Reorganizing the Conditions of Human Life

4.1 The first opportunity: reducing the human labor necessary to sustain basic life

The deepest civilizational significance of automation is not how many fewer people firms employ. It is whether the time and cognitive labor humanity needs to maintain its material existence can decline sustainably. If AI and automation raise productivity, lower the costs of education and public services, and make these gains widely accessible, humanity may for the first time reduce, at a larger scale, the structural pressure of “having to continually sell large amounts of one’s time to secure basic subsistence.”

This does not mean that work will disappear or cease to be meaningful. It means that work can be partially loosened from subsistence-driven compulsion, giving people better conditions for choosing labor, care, creation, and public participation worthy of long-term commitment.

4.2 The second opportunity: broadening access to capabilities once reserved for large institutions and a few experts

Generative AI and general-purpose tools can lower barriers to translation, coding, design, analysis, organizing materials, and professional communication. UNDP’s 2025 report emphasizes directing AI development toward human capabilities, choices, and complementarity [2]. If this capability expansion reaches beyond the wealthiest institutions into small firms, schools, communities, developing regions, and ordinary individuals, technological progress could broaden the social boundary of “who can do complex things.”

4.3 The third opportunity: stronger capabilities for addressing science, medicine, and planetary problems

AI has already entered drug, protein, weather, materials, astronomy, and biological research, expanding search and simulation spaces for some scientific tasks [1][19]. The WHO likewise treats multimodal models as general-purpose capabilities potentially applicable to clinical care, patient services, education, research, and drug development, while emphasizing errors, biases, automation bias, and governance risks [15].

If these capabilities are used to address disease, climate, energy, agriculture, disasters, and basic science, humanity may solve more rapidly problems previously constrained by information and computational complexity. This opportunity belongs to more than the technology industry: it could directly change healthy life expectancy, environmental carrying capacity, and global development.

4.4 The fourth opportunity: education can finally move further from knowledge scarcity toward human development

As high-quality explanations and practice become cheaper, education can invest more scarce resources in processes that are harder to replace through standardization: asking questions, deep understanding, practice, relationships, cooperation, embodied experience, ethical judgment, creation, and public participation. Vygotsky and the situated-learning tradition have long explained that capabilities form through social interaction and progressive participation [27][28]. AI can expand this process, but “giving answers” must not be mistaken for “human development.”

4.5 The fifth opportunity: one person’s capabilities and care can reach more real people

AI can cross linguistic, temporal, and professional barriers, extending a teacher’s knowledge, a doctor’s health education, a researcher’s explanations, or an ordinary person’s care to more people. What is most worth pursuing here is not “simulating someone to exist in their place,” but extending the capabilities, experience, and goodwill of flesh-and-blood people under explicit authorization and traceable responsibility.

4.6 The sixth opportunity: humanity is forced to ask anew what development is ultimately for

As doing things faster, more cheaply, and with greater output becomes easier, efficiency itself becomes increasingly inadequate as civilization’s ultimate answer. A civilization that can only keep expanding production and optimization, without knowing why human time should be freed, why relationships and choices should be protected, or why broader participation in public life matters, may possess powerful systems yet lack direction.

AGI’s greatest opportunity may be not to complete the world on humanity’s behalf, but to force humanity, for the first time, to decide seriously what we want civilization to be for once survival and production no longer consume all our capabilities.

5 | What Cannot Happen Automatically: Technological Progress Does Not Produce a Good Civilization by Itself

The error technological optimism most readily makes is not overestimating a particular model score, but translating “possible capabilities” directly into “civilizational outcomes.” A paper on civilizational direction must clarify the transformations in between.

Technological change	Common automatic inference	The transformation still required
More capable AI	Everyone naturally becomes more capable	Access, education, organizational design, and retention of human capabilities
Higher productivity	Everyone naturally becomes freer	Ownership, distribution of gains, working time, basic protections, and bargaining structures
More automation	People naturally have more time	Whether time saved is returned to people or filled again by new performance demands

Technological change	Common automatic inference	The transformation still required
More abundant information	Society naturally becomes more rational	Judgment of evidence, media ecosystems, education, and capacity for public discussion
Easier connection	Relationships naturally improve	Real reciprocity, responsibility, presence, exit, and freedom from manipulation
More accurate prediction	Governance naturally becomes more legitimate	Authorization, explanation, appeal, responsibility, rights, and democratic procedures
More stable systems	Civilization naturally improves	The position of subjects, basic rights, exit, and whether the direction is worthwhile
More intelligent AI	AI naturally has greater authority to decide	Capabilities and legal / political standing must be justified separately

The first principle of Civilization Leap is therefore neither “trust technology” nor “stop technology,” but reject automatic conversion. Every technological dividend requires questions: who actually receives it; whether human capabilities grow alongside it; whether power is being silently transferred; whether exit remains available; and whether life is genuinely improving.

The OECD’s AI Principles, updated in 2024, include inclusive growth, human capability enhancement, human rights, democratic values, safety, the ability to stop systems, and accountability within AI policy [12]. UNESCO’s 2021 Recommendation similarly treats human rights, dignity, human oversight, the environment, education, and culture as shared boundaries [13]. The EU AI Act establishes a legally binding governance framework through risk tiers, specific obligations, and requirements for high-risk systems [37]. These international frameworks are not equivalent to this paper’s civilizational goals, but they show that real-world governance can no longer use the “pace of innovation” as its sole measure.

6 | Direction: The Baseline of Continued Existence, Freedom from Survival Insecurity, Happiness in Life, and a More Meaningful Life

6.1 Continued existence is a baseline that must not be breached, but not civilization’s endpoint

Any civilizational direction must first preserve the conditions under which life and civilization can continue to exist. High-consequence risks, including AI militarization, biological risks, loss of control over critical infrastructure, and ecological and energy constraints, must not be deprioritized because a long-term vision is appealing. The International AI Safety Report 2026 places misuse, malfunction, and systemic risks alongside one another and notes that existing technical safeguards remain incompletely reliable [16].

Yet continued existence answers only “Can we carry on?” A highly stable system that leaves the vast majority of people in enduring scarcity, passivity, displacement, and lack of choice is still not the fully realized civilization this paper seeks.

6.2 The first goal: freedom from survival insecurity

Freedom from survival insecurity does not mean that “no one will ever face difficulties.” It means that access to basic food, shelter, basic health care, education, safety, and the conditions for continuing to live becomes progressively less liable to collapse altogether because of a single episode of unemployment, illness,

technological displacement, or circumstances of birth. In particular, basic human subsistence should not depend indefinitely on “first proving sufficient market value” as its only prerequisite.

The AGI era makes this goal more realistic because the marginal costs of production, knowledge, and services may fall. It does not make the goal automatic because productivity gains, ownership, public finances, and resource allocation remain institutionally determined.

6.3 The second goal: happiness in life

After staying alive, people still need real lives. “Happiness in life” here is not a psychological score that must be measured uniformly. It means having more self-directed time, health, safety, relationships, freedom, enjoyment, culture, nature, choice, and capacity to exit. High consumption or high platform engagement cannot substitute for happiness.

Free time does not automatically produce happiness, either. If the time automation releases is merely filled by new attention platforms, performance demands, or isolation, civilization has not genuinely returned capability to people. The real direction is to give people both time and practical conditions to love, care, learn, play, create, rest, participate in public life, and live together.

6.4 The third goal: a more meaningful life

“A more meaningful life” does not mean civilization assigning everyone a uniform answer. Sources of meaning can differ: relationships, creation, care, knowledge, responsibility, public participation, faith, exploration, experience, and existence itself can all generate meaning. What civilization can do is not prescribe why you should live, but expand the conditions under which you can genuinely enter these sources of meaning.

Civilization Leap does not assign meaning to people; it enables more people to have the conditions for generating meaning.

6.5 The three goals form a progression, but cannot be exchanged for one another

Freedom from survival insecurity → happiness in life → a more meaningful life is a directional progression, not three interchangeable KPIs. Higher overall safety cannot justify continually sacrificing some people’s basic rights; an increase in average happiness cannot justify eliminating people’s choices; and “greater meaning” cannot justify demanding that individuals accept deprivation and coercion.

These three goals acquire renewed practical weight in the AGI era because technology is simultaneously loosening the conditions of subsistence, labor, knowledge, and time. In the past, many societies first had to devote substantial civilizational capabilities to maintaining basic production. If necessary labor and knowledge scarcity decline in the future, humanity has reason to extend civilization’s goals further toward life and meaning.

7 | The Human Place: Civilization Leap Means Real Lives and Diverse Sources of Meaning, Not Merely “Something Left to Do”

7.1 “Finding tasks for people” is not a civilizational answer

The most immediate response to automation is often “How can we make sure people still have jobs?” Work certainly matters. But if civilization merely invents low-value tasks to prove that people remain useful, it locks human value back into “whether the system can use them.” The real issue is not whether systems still have a use for people, but whether people retain a genuine place in judgment, responsibility, development, relationships, creation, and the formation of shared direction.

7.2 Structural meaning: the mid-level interface through which meaning enters real life

This paper uses “structural meaning” in the mid-level sense developed in Meaning Dynamics: it concerns how capabilities, needs, and connections form sustainable linkages for action through relationships, pathways, roles, contributions, participation, and institutions. It does not answer ontological questions about meaning or reduce private experience to structural variables. Its function is simply to make visible that although institutions cannot define the meaning of someone’s life, civilization can expand or constrict their opportunities to enter real relationships, assume responsibility, contribute, develop, and be needed.

7.3 From a single occupational center to multiple sources of meaning

Occupations carry too many functions in modern civilization: income, identity, social connection, time, development, contribution, and recognition of value are often bundled into one job. When automation disrupts occupations, the problem may therefore appear simultaneously as economic pressure, destabilized identity, fewer relationships, and shrinking sources of meaning. A more robust future direction is not to make “work” disappear, but to ensure that a person’s world of meaning no longer depends exclusively on one occupational path.

Source of meaning	Possible real-world entry points	Contraction to avoid in the AGI era
Relationships and love	Family, partners, friends, communities	Fully platformizing relationships; replacing companionship with simulations that bear no responsibility
Creation	Art, design, writing, invention, crafts, entrepreneurship	Leaving people only final-stage touch-ups and removing the processes of making work and forming judgments
Knowledge and development	Learning, research, mentorship, exploration	Replacing understanding with answers and long-term capability formation with systems
Care and being needed	Childcare, eldercare, medicine, neighborhood support, volunteering	Treating care only as an inefficient cost while ignoring relationships and responsibility
Work and contribution	Occupations, projects, maintenance, services, production	Fragmenting jobs into shallow tasks so contributions cannot consolidate into roles and earned trust
Public participation	Civic discussion, communities, governance, public services	Allowing opaque systems and a few interfaces to monopolize public decisions
Faith and spiritual life	Religion, traditions, remembrance, ultimate concerns	Allowing AI’s functional roles to impersonate representatives of traditions, clergy, or deceased people themselves
Exploration and existence	Travel, nature, embodied experience, play, the unknown	Filling free time again with consumption and algorithmically captured attention

The more diverse the sources of structural meaning, the less likely an entire life is to collapse when one pathway becomes unstable. Yet diversity does not mean that more pathways are always better. Without time, resources, relationships, and institutional support, a multitude of “choices” merely becomes a new overload. What civilization should genuinely expand is access to sustainable pathways through which responsibility and development can take shape.

8 | Shared Problems: How Disciplines Can Address Civilizational Events Together Without Erasing Professional Boundaries

Civilizational events in the AGI era matter because they simultaneously change the objects of multiple disciplines. Disciplines engage not to prove that “Civilization Leap theory is correct,” but because real-world problems have already crossed boundaries within which they could previously be handled independently.

Shared event	Disciplines able to assess it first	Disciplines that must re-engage	Civilizational question that must not be omitted
AI restructures labor	AI, economics, organizational studies	Education, sociology, law and politics, philosophy	After productivity rises, how are basic life, identity, development, and being needed reorganized?
AI enters education	Education, cognitive science, AI	Ethics and law, society and culture, economics, politics	As answers become cheaper, how do people still develop judgment, relationships, civic capabilities, and a self?
AI enters medicine and decisions about life	Medicine / life sciences, AI	Ethics and law, economics, sociology, philosophy	After accuracy improves, who decides, who is responsible, and who receives care?
AI enters government	Politics, law, public administration, AI	Epistemology, economics, sociology, systems science	Does more accurate prediction create de facto power? Can citizens still understand, question, and appeal?
AI enters relationships and culture	Sociology and anthropology, cognitive science, AI	Religion, ethics and law, economics, philosophy	Does improved access preserve real relationships, boundaries of representation, and cultural diversity?
AI militarization	AI, strategic security / international relations	Systems science, ethics and law, epistemology, economics	Do locally rational security choices aggregate into shorter correction windows and risks to shared continued existence?
Computing and resource expansion	Energy / Earth systems, engineering, economics	Politics, international relations, sociology	How does AI growth affect electricity, materials, water, space, and Earth's carrying capacity?

8.1 A good metric in one domain may not be a good metric in another

Efficiency may be important in computer science and business management, but in education it cannot replace understanding and development; in relationships and mourning, it may not even be a primary goal. Accuracy matters in medicine, but cannot determine patient rights and responsibility by itself. Prediction is valuable in governance, but cannot replace legitimacy and appeals. Moving beyond disciplinary confines is not about making all disciplines use one metric: it is about knowing what further judgments a discipline’s version of “good” must face when it enters another domain.

8.2 Aggregation matters more than simple “integration”

Civilizational risks often arise not from an application alone, but from many locally rational choices simultaneously adopting similar infrastructure and optimization logics. AI in education, medicine, office work, government, and research may each improve efficiency. Yet if they all rely on a few models, cloud platforms, and data channels, their aggregation may produce shared dependence, correlated failure, de facto power, and difficulty exiting. This is a new object of inquiry requiring systems science, economics, politics, and engineering to work together.

Similar aggregation may occur between culture and models. Experiments on specific creative tasks have found that generative AI can improve individual creative performance while reducing the collective diversity of content [21]. Other research shows that, under particular recursive training conditions, models continually learning from model-generated data may undergo distributional degradation or even model collapse [22]. Neither finding directly proves that a complete feedback loop has already formed in the real world. Together, however, they justify a refutable category C question: if people increasingly obtain their forms of expression from similar models, and models increasingly learn from those expressions, might the space of cultural and cognitive exploration narrow on both sides?

Disciplines pursue truth separately; civilization bears the consequences together.

9 | Method One: Start with Real Problems and Genuinely Affected People

9.1 Do not start with a complete blueprint

Civilization Leap does not begin by drawing a perfect society and then seeking reality to validate it. A workable method starts from problems already occurring: a type of work being restructured, learning dependence emerging in a school, de facto algorithmic power appearing in a public system, a community lacking care, or a resource bottleneck constraining AI expansion. The more concrete the problem and the more real the affected people, the more testable the professional judgments and consequences.

9.2 Genuinely affected people are not “user personas”

Who bears risks, receives benefits, loses pathways, or cannot exit must not be left to designers’ imagination. Both the WHO’s work on medical AI and UNESCO’s guidance on educational AI emphasize stakeholder participation, specific use contexts, and human oversight [14][15]. A civilizational method must turn patients, students, teachers, workers, families, disadvantaged groups, communities, and future generations from abstract “beneficiaries” into people who can change the proposal.

9.3 Address the most concrete difficulties first, then discuss generalizability

A mechanism working in one city, school, or organization does not make it a universal civilizational answer. Conversely, a small-scale failure need not invalidate the direction itself. Real-world trials turn abstract causation into observable consequences: who actually benefits, which costs have been overlooked, what conflicts of rights emerge, whether new dependencies arise, and whether exit is genuinely feasible.

10 | Method Two: Allow Old and New Arrangements to Coexist, Compare and Validate, and Permit Failure and Exit

10.1 A new civilization does not mean “declaring the old world over”

Existing institutions sustain rights, livelihoods, credit, assets, family responsibilities, and social continuity. Even if a new mechanism is theoretically better, a vision alone cannot justify demanding everyone’s immediate migration. A more reasonable path allows old and new mechanisms to coexist within tolerable bounds, lets people compare their actual effects, and permits a better reality to earn trust gradually.

10.2 “First show that it is better” matters more than “first prove that it is correct”

Civilizational direction cannot be fully proven on paper. What a new resource, educational, collaborative, or public mechanism must actually demonstrate is whether, at the same or lower cost, it gives more people better living conditions, reduces coercion and irreversible harm, and preserves more choice and scope for correction.

10.3 The right to fail is infrastructure for civilizational innovation

If every failed experiment means permanent loss of standing, devastating losses in daily life, or an irreversible penalty to one’s status, people will not genuinely try new pathways. Civilization Leap must lower the cost of controllable failure, giving people and organizations opportunities to exit, start again, revise, and return to previous pathways.

10.4 Exit is a real capability, not a product button

Genuine exit means alternative services remain available after use stops; data and identity are not permanently locked in; basic living conditions and rights do not suffer because of exit; and organizations retain manual takeover and backup capabilities. The OECD AI Principles emphasize mechanisms for override, repair, and safe decommissioning when systems may cause undue harm or exhibit undesired behavior [12]. NIST’s risk frameworks also emphasize lifecycle governance [17][18]. At the civilizational level, we must further ask whether individuals, organizations, and states genuinely have alternatives, rather than merely nominal permission to leave.

11 | Method Three: Use AI to Expand Human Capabilities While Retaining Judgment, Responsibility, Development, and Direction

11.1 A human in the loop is not necessarily a human in the position of a subject

A system can retain a “human confirmation” button while handing actual problem definition, evidence organization, comparison of options, and judgments of responsibility to AI in advance. The human remains in the workflow, but only as the final signatory. Research on automation has long shown that excessive reliance can produce skill degradation, supervisory burdens, and difficulty taking over [33][34].

11.2 Judgment develops through real acts of judgment

Human judgment, creativity, and responsibility are not static assets stored in the brain. They develop through sustained engagement with real problems, mistakes, explanation, consequences, and collaboration. If education and organizations eliminate these processes for efficiency, even better short-term results may reduce the future reproduction of robust human subjecthood.

11.3 The minimum direction for AI augmentation

Human position to preserve	What AI can appropriately do	What must not occur silently
Understanding	Organize information, explain different materials, and provide comparisons	System summaries becoming the sole entry point to facts
Judgment	Offer candidate options, risks, and counterexamples	Default recommendations becoming de facto decisions
Creation	Expand materials, simulations, prototypes, and expressive capabilities	Keeping people permanently in a post-editing role
Responsibility	Record processes, flag risks, and support tracing	Responsibility drifting among model, platform, and organization
Development	Personalized feedback, practice, and assistance to mentors	Eliminating difficulty, failure, mentoring, and independent capability formation
Relationships	Translation, reminders, and extension of capabilities	Replacing real reciprocity and presence with simulations that bear no responsibility
Direction	Provide scenario analysis, comparisons of options, and visualization of consequences	Superior capability automatically gaining final decision-making authority

11.4 Organizational incentives must account for long-term capacities for subjecthood

Organizations usually find it easier to measure the work hours, costs, and delivery time AI saves than the mentoring, judgment training, experiential redundancy, and long-term capacity for responsibility it removes. If these slow-moving variables are persistently excluded from organizational decisions, replacement of human subjecthood will enjoy a structural short-term advantage. Civilizational direction does not prohibit automation; it requires organizations to treat future takeover capability, learning, responsibility, and human development as real assets too.

12 | Method Four: Change Priorities by Phase Without Making Emergency Arrangements Permanent

12.1 The launch phase: first avoid losing the future and the window for correction

When AI capabilities expand rapidly, strategic competition accelerates, and institutions remain immature, the launch phase's primary task is to preserve shared continued existence, basic rights, and space for further experimentation. Research on complex-system safety has long warned that high-consequence accidents may arise from interactions among components, organizational processes, and control structures, rather than a single-point error [31][32]. High-consequence systems therefore need stronger monitoring, responsibility, stopping, and takeover capabilities; societies must avoid large-scale irreversible displacement caused by one technological wave; and new mechanisms must quickly demonstrate whether they can produce real benefits.

12.2 The transition phase: make old and new structures comparable in practice

The important task of transition is not to expand temporary power, but to move new mechanisms beyond a few trials into an institutional environment where results are replicable, comparable, and open to appeal. At this stage, exceptional launch-phase arrangements must begin to be dismantled, preventing temporary permissions granted “for safety and speed” from hardening into permanent power.

12.3 The stable phase: restore rights, exit, balance, and quality of life to the center

Once basic continued existence and the new operating conditions are relatively stable, civilization should no longer govern everyday life through the launch phase’s emergency logic. The stable phase places greater emphasis on regular boundaries of rights, balances of power, long-term governance, procedures, reversibility, and quality of life. The more efficient technology becomes, the more important it is to avoid subjecting all of life to unified optimization again.

12.4 Self-maintenance means society’s continuing capacity for correction, not “a system running automatically”

A mature civilization does not need a central designer who is always right. It needs professional research, public discussion, institutional correction, engineering maintenance, and real-life experience to enter continually. Ostrom’s research on polycentric governance shows that complex public affairs admit institutional possibilities beyond a single central controller [24]. This does not directly establish any particular governance structure for a future civilization, but it supports keeping polycentric experimentation, comparison, and correction open as directions. The core of self-maintenance is not locking institutions in place, but enabling them to discover that they are wrong and to change.

13 | Goal Realization: Observing Freedom from Survival Insecurity Becoming a Reality

Freedom from survival insecurity cannot be reduced to an income threshold or replaced by the name of a welfare system. It is a set of real-world directions: basic life becomes progressively less liable to collapse altogether because of market fluctuations, technological displacement, illness, family background, or one failure; and people have sufficient choice that maintaining subsistence does not require accepting any job, any platform terms, or any technological dependence.

Direction to observe	Signs of movement toward the goal	False realization to guard against
Access to basic resources	More stable access to food, shelter, basic health care, education, and basic digital connectivity	Nominal guarantees whose applications are complex, punitive, or practically inaccessible
Resilience to employment shocks	Maintaining life, learning, and entry into other pathways after technological displacement	Short-term subsidies concealing lasting displacement and lock-in of eligibility status
Basic rights do not vary with scores	Exiting platforms or refusing certain AI uses does not affect basic subsistence or civil rights	Contribution, credit, or behavior scores becoming thresholds for basic resources
Recoverability from failure	One failed venture, transition, or learning attempt does not permanently exclude a person from social structures	Failure causing long-term debt and blocked status and pathways

Direction to observe	Signs of movement toward the goal	False realization to guard against
Narrowing regional and group disparities	AI capabilities and basic services expand into peripheral regions and disadvantaged groups	Average indicators rise while technological gains accrue only to a few centers

“Freedom from survival insecurity” does not eliminate individual effort, markets, or differences, nor does it promise unlimited resources. It changes only one civilizational baseline: basic subsistence should not be founded indefinitely on “first having to prove oneself useful enough.”

14 | Goal Realization: Observing Happiness in Life Becoming a Reality

Happiness in life is far more complex than subsistence and cannot be replaced by a uniform satisfaction measure. Sen’s capability approach reminds us to evaluate development not only through resources and output, but also through what people can actually do and become [23]. Self-determination theory has long emphasized the connections of autonomy, competence, and relatedness with motivation and well-being [25]. Research on social relationships likewise continues to show the importance of human connection for health and long-term well-being [26]. These studies do not directly prove this paper’s civilizational goals, but they demonstrate that quality of life cannot be understood through GDP or efficiency alone.

Dimension of life	Directional observation question
Self-directed time	After necessary labor and administrative burdens decline, do people actually have more time they can decide how to use?
Health and safety	Does technology improve access to prevention, diagnosis, treatment, and care without introducing new inexplicable risks?
Real relationships	Do people have more time and space for family, friendship, community, and lasting mutual support, rather than platform interactions alone?
Choice and exit	Can people refuse a particular job, platform, algorithmic recommendation, or organizational arrangement without incurring unbearable costs?
Culture and enjoyment	Do cities, communities, education, and digital spaces make room for play, art, nature, exercise, and non-instrumental living?
Freedom from continual optimization	Does life allow slowness, leisure, chance, failure, privacy, and existence that produces no metric?

AGI may improve happiness in life, or it may carry the logic of efficiency further into homes and private life. The criterion for civilizational direction is not “more technology use is better,” but whether technology ultimately expands people’s real choices and the richness of their lives.

15 | Goal Realization: Observing the Expansion of Conditions for a More Meaningful Life

15.1 We do not measure “how much meaning you have”

Meaning is particularly vulnerable to misuse by governance systems. Once “meaningfulness” becomes a uniform score, civilization begins prescribing which lives are worthwhile. This paper rejects that direction. What can be observed is not how much a person’s “meaning is worth,” but whether civilization expands the conditions for generating meaning.

We do not measure “how much meaning you have”; we observe whether civilization is expanding people’s conditions for generating meaning.

15.2 Directional observation of the conditions for meaning

Condition	Observable real-world change	What it must not be reduced to
Real relationships	Care, friendship, family, and community have time, space, and reciprocity	Number of connections or frequency of messages
Being needed and contributing	Real pathways for diverse contributions such as maintenance, care, creation, bridge-building, and public service	A uniform total contribution score
Development	People can move from shallow tasks to substantial roles, continually accumulating judgment, responsibility, and capability	Number of credentials or platform levels
Creation and exploration	Low-barrier opportunities to enter art, science, projects, nature, and the unknown	Volume of content or traffic
Public participation	More people can understand, question, and influence shared life	Passive likes or counts of formal votes
Spiritual life and ultimate concerns	Different religions, traditions, and personal answers have autonomous space and real community support	AI deciding individuals’ beliefs or simulating sacred authority
Life space not predefined	People can still change pathways, reject a prescribed optimum, encounter chance, and start again	Algorithms preplanning a “best life” for everyone

15.3 Sources of meaning must take culturally situated forms, not a globally uniform one

East Asian societies may place greater weight on relationships and responsibility; some religious communities closely connect daily life to a transcendent order; modern liberal societies may emphasize individual choice and civil rights; local communities may rely on informal mutual aid and oral traditions. Shared questions of structural meaning can be compared, but the forms that sustain it are necessarily culturally situated.

Civilization Leap therefore cannot turn one model of family, occupation, religion, or public participation into a uniform template for humanity. The genuinely shared direction is simply to enable different people to have more

real, sustainable, non-coercive pathways for forming their own relationships, responsibilities, creations, development, and answers.

16 | From Direction to the Lifeworld: The Responsibilities of Vision, Institutions, Engineering, and Real-World Systems

16.1 Thought provides guidance on “why” and “where”

The paper’s main responsibility has now been fulfilled: to explain why AGI opens a window for civilizational choice; why technological dividends do not automatically become human dividends; why the three goals form a direction worth debating and testing; and how open methods can move us toward it gradually. Engineering cannot substitute for thought.

16.2 Living Civilization lets readers see “what it would be like to live there”

Academic argument can explain direction and causation, but concepts alone make it difficult to judge whether a future would be worth living in. The Living Civilization trilogy gives everyday form to a future world: how work, family, education, relationships, communities, public life, and technology might be reorganized. It is a vision and an entry point for further reading, not evidence for this paper’s empirical claims.

16.3 Spark, Nebula, AGI-COS, Human-COS, and others are only candidate implementations

Institutional and engineering systems address “how to give this practical form.” Spark here is a candidate system for extending individual capabilities and care. Nebula is a candidate framework for civilizational collaboration and operation. AGI-COS is a candidate engineering-governance framework for machine-side capabilities, permissions, auditing, stopping, and exit. Human-COS is a candidate system supporting multi-actor, multi-perspective exploration of consequences and judgment before high-consequence choices become irreversible. They remain on a research and validation track: they may provide support, undergo trials, fail, or be exited. Membership in the Civilization Leap framework cannot automatically make them correct, still less make them evidence for this paper’s directional judgments.

16.4 The causal order from direction to engineering must not be reversed

Thought sets direction; disciplines pursue truth; institutions set boundaries; engineering gives reality practical form. Reality, in turn, revises them all.

If readers must first learn many internal identifiers to understand why a mechanism exists, the order of explanation is usually wrong. When “allow people to exit,” “return responsibility to real subjects,” or “let new ways of living first prove better” expresses the point clearly, complex protocols should not replace thought at the outset.

17 | How Reality Can Refute Us: Failure, Counterexamples, and Civilizational Self-Correction

The most dangerous condition for a paper on civilizational direction is treating “worth pursuing” as “therefore certain to work.” This paper must specify explicit ways for contrary evidence to force revision.

This paper's judgment	What evidence would compel revision
AI may constrict human judgment and development over the long term	If longitudinal, multi-context studies consistently show that independent human capability, takeover capability, and capacity for responsibility generally improve after AI is withdrawn, the assessed risk of "thinning subjecthood" should be reduced
Productivity dividends do not automatically become freedom	If AI productivity gains across different institutional settings generally and persistently reduce working hours and improve income security without widening distributional disparities, concerns about distributional mediation should be narrowed
The single occupational center will loosen	If occupations still stably sustain identity, relationships, development, and meaning after large-scale automation, while non-occupational pathways fail to provide lasting support, the judgment about reorganization of meaning sources should be revised
Infrastructure concentration will produce systemic dependence	If openness, competition, interoperability, and public infrastructure preserve genuine substitutability and independent exit over the long term, the judgment about aggregate dependence should be narrowed
The three goals have substantial cross-cultural appeal	If different civilizations, groups, and sustained public discussion repeatedly show that the three goals omit irreplaceable civilizational goals, the author should reopen the normative framework rather than explain dissent as "not yet understanding the theory"
Coexistence of old and new arrangements and the possibility of exit are preferable methods	If reality repeatedly shows that coexistence significantly increases irreversible risk under particular high-consequence conditions, or that exit mechanisms cause greater harm, the methods should be reset by domain and phase

17.1 Self-correction is not a fourth goal

Continued existence, freedom from survival insecurity, happiness in life, and a more meaningful life describe the states civilization seeks. Self-correction describes how civilization avoids becoming trapped by its own earlier answers. It is an operating capability and should not be mechanically added as a fourth goal.

17.2 Direction must also remain open to change through different civilizational experiences

A normative commitment cannot be "falsified" by a single observation in the manner of an experimental hypothesis. It can, however, be forced into revision by stronger arguments, long-term consequences, affected people's experiences, and cross-cultural counterexamples. If a supposedly "meaningful civilization" repeatedly requires coercion, if "governance for happiness" continually erodes freedom, or if "survival security" permanently eliminates exit, it is not the public that has failed: our formulation of direction or its institutional realization has.

17.3 The most valuable participation is not agreement

The most valuable participation is not agreement with this paper, but professional work capable of changing it.

Researchers can use longitudinal data to test capability enhancement and the thinning of subjecthood; economists can compare distributional institutions; educators can measure independent capability after AI is withdrawn; sociologists and anthropologists can supply counterexamples concerning how different civilizations sustain meaning; strategic researchers can test AI races and crisis stability; ecological and energy research can

recalculate AI expansion's physical limits; and engineers can demonstrate whether particular exit and redundancy designs are genuinely feasible.

Conclusion | Technology Opens Possibilities, Thought Determines Direction, and Reality Tests the Path

What truly changes humanity in the AGI era is not simply whether machines become more intelligent, but the technological extension of ever more knowledge, productive, analytical, organizational, and execution capabilities previously dependent on people. These growing capabilities will naturally enter firms, schools, hospitals, governments, families, and interstate competition. They need not wait for a day when the whole world jointly declares that "AGI has arrived."

This is also why civilizational direction cannot be postponed. Technological and organizational inertia already make choices: where efficiency matters, who bears costs, who receives judgment, whether time is returned to people, who controls infrastructure, what education cultivates, and what relationships become. If humanity does not actively answer, many civilizational directions will still be chosen, but the choosers will be market incentives, platform defaults, organizational KPIs, and competitive pressures.

This paper openly takes responsibility for a different direction. First preserve continued existence, and pursue freedom from survival insecurity as technological capability can reduce scarcity. Once subsistence is better supported, continue directing civilizational capabilities toward happiness in life that is real, free, rich, relational, and enjoyable. Then go further: do not prescribe meaning for people, but expand the conditions under which relationships, creation, care, learning, responsibility, public participation, faith, exploration, and existence itself can genuinely enter their lives.

This is neither a forecast nor a demand that all civilizations accept the same institutions. It proposes only a directional judgment that facts, professional expertise, different cultures, and real-world failures can continually change.

AGI's greatest significance is not how humanlike machines may ultimately become. It is that as machines take on more and more things that once required people, humanity must finally ask anew: what do we want civilization to exist for?

Technology opens possibilities, thought determines direction, and reality tests the path.

Appendix A | Civilizational Event-Discipline-Direction Matrix

Civilizational event	Main disciplines engaging together	What may be missed without further tracing	Question of civilizational direction
Widespread intelligent capabilities	AI, economics, education, organizational studies, sociology	Capability access and capability concentration coexist	Have more people genuinely become more capable?
Labor restructuring	Economics, organizational studies, education, sociology, law and politics, philosophy	Counting jobs alone misses development and being needed within jobs	Can productivity become time, security, and choice?
AI-based knowledge production	Epistemology, sciences, AI, education, political economy	Counting outputs alone misses truth-seeking, infrastructure, and authority over direction	As knowledge becomes faster, who can still judge independently?
AI-based education	Education, cognitive science, AI, ethics and law, society and politics	Looking only at grades misses relationships, development of subjecthood, and civic capabilities	As answers become cheaper, what should people develop into?
AI-based medicine	Medicine, life sciences, AI, ethics and law, society and economics	Looking only at accuracy misses decisions about life, care, and responsibility	As capabilities improve, who decides, who is responsible, and who receives care?
AI-based public governance	Politics, law, public administration, AI, epistemology	Looking only at administrative efficiency misses legitimacy, appeal, and de facto power	As government predicts better, how much capacity to understand and question do citizens retain?
AI-based relationships and culture	Sociology and anthropology, cognitive science, religion and culture, ethics, AI	Looking only at interaction and satisfaction misses real relationships, standing to represent, and cultural diversity	Should technology amplify relationships or replace them?
AI militarization	Strategic security, international relations, AI, systems science, ethics and law	Looking only at one side's capabilities misses races, misjudgment, and shared continued existence	How can states' local security avoid aggregating into shared insecurity?
Computing / energy expansion	Energy, Earth systems, economics, engineering, politics	Looking only at model capabilities misses physical carrying capacity and resource distribution	How can intelligence expand without exceeding Earth's and society's carrying capacities?

This matrix does not constitute a new disciplinary catalogue. The earlier joint review working backward from real-world pressures established that three areas of factual responsibility require independent professional engagement. The professional truth-seeking layer has therefore expanded from thirteen papers to sixteen: international relations and strategic security cover deterrence, crisis stability, and strategic interaction; ecology, Earth systems, and resources and environment cover energy, water, materials, land, and material carrying capacity; and medicine, public health, and health systems cover clinical effectiveness, patient safety, and population health. “First, second, and third pressures” refer only to an ordering of real-world irreversible pressures, not disciplinary value or manuscript quality, and no longer determine whether additional papers should be added.

Appendix B | Already Visible, Highly Probable Continuation, and Genuinely Open: A Checklist for Judgments About the Future

Layer	Writing rule	Typical examples	Incorrect formulation
Already visible	Specify time, object, and source; permit direct external verification	Expanding organizational AI adoption, exposure of job tasks, rising data center electricity consumption	“All industries have already been restructured by AGI”
Highly probable continuation	Specify conditions and routes for refutation; usually category C structural judgments	Work restructuring, organizational dependence, changing educational purposes, reorganization of meaning sources	“Mass unemployment is inevitable”
Genuinely open	Explicitly acknowledge what is unknown; do not make it a necessary premise for present institutions	AGI dates, machine consciousness, superintelligence, extreme abundance, mind uploading	“Superintelligence is certain to arrive, so today we must...”

Minimum pre-release check: if all assumptions about superintelligence, machine consciousness, and specific AGI arrival years were removed, the paper’s central directional argument should still hold. If it does not, the paper has mistakenly grounded civilizational choice in an open future.

Further Reading and Joint Participation | Letting Reality Continue to Change the Theory, Not Merely Finishing It

The paper has now explained why the AGI era opens a window for civilizational choice; which realities and highly probable changes deserve engagement today; why technological opportunities do not automatically become human dividends; which three-goal direction it chooses; how open methods that allow testing, exit, and revision can approach that direction gradually; and what facts could force it to reopen its conclusions. What follows proceeds by question rather than stacking documents in file order.

Upward | Why do these local changes ultimately become civilizational questions?

Turn to [Reassessing the Disciplines in the AGI Era: Disciplinary Boundaries, Interactions, and a Shared Research Agenda V2.4](#), and the five-volume Civilization Leap series. The former addresses how disciplines reconnect causal chains; the latter further develops why civilization needs to change and how institutions and the human place can be understood anew.

Public entry to the five-volume Civilization Leap series: <https://www.civitas.top/en/documents/p02-five-001/>

Downward | How can these directions enter verifiable institutions, engineering, and real-world systems?

Continue with Spark, Nebula, AGI-COS, Human-COS, and the other institutional and experimental papers described in Section 16.3. They are candidate implementations, not proof of this paper's direction. Their real value is in subjecting that direction to implementation, failure, exit, and real-world consequences.

[Spark](#)

[Nebula](#)

[AGI-COS](#)

[Civilization Leap Library](#)

Leftward | How does reality actually propagate into these civilizational judgments?

Return to the first paper, [AI Application Prospects and Risks in the AGI Era: From Capability Expansion to Civilizational Consequences V1.2](#), and the corresponding in-depth disciplinary papers. The first paper addresses real-world applications, organizational mediation, cross-domain feedback, and aggregate consequences. The sixteen disciplinary papers are responsible for their own facts, theories, limits of evidence, and competing explanations. This paper cannot seek truth on their behalf.

Rightward | If you disagree, which contrary directions most deserve investigation?

Challenge the paper directly through these questions: Could AI improve independent human capability more reliably than this paper anticipates? Could productivity gains diffuse naturally under existing institutions? Will occupations remain the center of meaning over the long term? Could competition and open source sufficiently address infrastructure concentration? Do the three goals omit more important civilizational goals? Could "exit and coexistence of old and new arrangements" increase risk in some high-consequence domains? Any counterexample that improves the factual quality of these questions matters more than conceptual agreement.

Forward | From reading to research, institutions, engineering, and shared creation

Researchers can build long-term cross-domain datasets; educators can compare independent capability after AI is withdrawn; economic and organizational researchers can trace where time savings and gains actually go; strategic-security researchers can model AI races and crisis stability; ecological and energy researchers can calibrate the real limits of computing expansion; engineers can test whether exit, takeover, redundancy, and auditability are genuinely feasible in highly capable systems; and members of the public and affected people can reveal consequences for daily life that the theory has missed.

The aim is not for more people to finish reading Civilization Leap, but for more people to help bring Civilization Leap into being.

The future pathway | What kind of world might people inhabit when these principles enter real life?

Turn to the Living Civilization trilogy. It translates an arguable direction back into an imaginable world of work, family, education, relationships, cities, and public life. Vision is not evidence, and readers may reject its future and propose alternative civilizational possibilities.

Living Civilization trilogy: an entry point to a vision of life in a future civilization

Meaning Dynamics: an entry point to the mid-level theory of structural meaning and the reorganization of sources of meaning

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

Generative AI Use Statement

Generative AI assistance was used in identifying leads for public-source searches, comparing existing manuscript versions, organizing structure, drafting portions of the text, checking references, arranging tables, and generating the layout. Generative AI outputs are not treated as empirical data, authoritative sources, independent evidence, or peer-review opinions. The research questions, civilizational direction, A/B/C/F argumentative statuses, selection of materials, verification of facts and citations, normative judgments, relationships within the theoretical framework, and final text were all reviewed and approved by the author, who is responsible for all content and its academic accountability. AI is not listed as an author. V1.1 makes controlled cross-manuscript consistency revisions to V1.0: it synchronizes the status of the sixteen disciplinary responsibilities, versions of the two synthesis papers, wording of the boundary around AI subjecthood, minimum explanations of candidate implementations, and stable public entry points. It does not change the main structure of three layers of judgments about the future, three goals, four methods, and observation of goal realization. Dynamic external facts were reviewed with particular attention through August 26, 2026. On formal submission or publication, tool names, versions, dates, and scope of use should be added as required by the target journal, publisher, or institution.

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] United Nations Development Programme. (2025). Human Development Report 2025: A Matter of Choice—People and Possibilities in the Age of AI. UNDP. <https://hdr.undp.org/content/human-development-report-2025>
- [3] 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>
- [4] 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>
- [5] 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>
- [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] 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>
- [8] 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>
- [9] International Energy Agency. (2025). Energy and AI. IEA. <https://www.iea.org/reports/energy-and-ai>
- [10] OECD. (2025). Competition in Artificial Intelligence Infrastructure. OECD Roundtables on Competition Policy Papers, No. 330. <https://doi.org/10.1787/623d1874-en>
- [11] United Nations. (2024). Global Digital Compact, Annex I to the Pact for the Future, A/RES/79/1. <https://www.un.org/pact-for-the-future/en/annex-i-global-digital-compact>
- [12] OECD. (2024). OECD AI Principles, updated 2024. OECD.AI. <https://www.oecd.org/en/topics/ai-principles.html>
- [13] UNESCO. (2021). Recommendation on the Ethics of Artificial Intelligence. Paris: UNESCO.
- [14] UNESCO. (2023). Guidance for Generative AI in Education and Research. Paris: UNESCO. Updated 2026.
- [15] 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.
- [16] International AI Safety Report. (2026). International AI Safety Report 2026. <https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026>
- [17] 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>
- [18] 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>
- [19] 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>
- [20] 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>
- [21] 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>
- [22] 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>
- [23] Sen, A. (1999). Development as Freedom. Alfred A. Knopf.

- [24] 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>
- [25] Ryan, R. M., & Deci, E. L. (2000). Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being. *American Psychologist*, 55(1), 68–78. <https://doi.org/10.1037/0003-066X.55.1.68>
- [26] Holt-Lunstad, J., Smith, T. B., & Layton, J. B. (2010). Social Relationships and Mortality Risk: A Meta-analytic Review. *PLoS Medicine*, 7(7), e1000316. <https://doi.org/10.1371/journal.pmed.1000316>
- [27] Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511815355>
- [28] Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Harvard University Press.
- [29] Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- [30] Kestin, G., Miller, K., Klaes, 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>
- [31] Perrow, C. (1984). *Normal Accidents: Living with High-Risk Technologies*. Basic Books.
- [32] Leveson, N. G. (2012). *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press. <https://doi.org/10.7551/mitpress/8179.001.0001>
- [33] 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)
- [34] Parasuraman, R., & Riley, V. (1997). Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
- [35] International Committee of the Red Cross. (2025). *Autonomous Weapon Systems and International Humanitarian Law: Selected Issues*. ICRC Position Paper.
- [36] OECD & UNESCO. (2024). *G7 Toolkit for Artificial Intelligence in the Public Sector*. OECD Publishing. <https://doi.org/10.1787/421c1244-en>
- [37] European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. *Official Journal of the European Union*.