

Civilization Leap Shared Material Notes and Further Reading

Companion to the R10.10 English Volumes

English Translation

Derived from R10.10

Zijun Fu

Civilization Leap Research Group

2026

About This Companion

This companion preserves the complete set of sixty-two material notes from Civilization Leap v1.21-R10.10, together with the upgraded public further-reading pages. It accompanies the separately translated volumes. It does not contain the five-volume main text.

Shared Material Notes

On the Historical Scenes, Cases, Future Extrapolations, and Literary Metaphors in This Book

The following notes collect the material referenced by light superscript markers in the main text. They identify the material, limits of sources, status of thought experiments, and literary metaphors; they do not carry the main argument. Unless delaying an explanation would materially mislead a general reader about facts, subjects, rights, institutional force, or present versus future status, qualifications of this kind are generally no longer inserted into scenes, strong propositions, sustained parallel passages, or closing resonances.

1 | Volume One, Chapter 1 | The status of “An Endgame in the Igloo”

The opening is a literary composite constructed from polar scarcity conditions and related ethnographic and travel accounts. It corresponds to no single verifiable person or event and does not generalize about Inuit people or any Arctic community. Its intellectual function is to reveal the structural question of how extreme scarcity may bring “usefulness” close to “whether one may remain.”

2 | Volume One, Chapter 1 | The status of ubasute material

Stories of abandoning older people, such as Japan’s ubasute, exist mainly in folklore, literature, and cultural memory. They cannot directly serve as empirical evidence of a widespread historical institution. The book uses them to present a long-standing cultural fear connecting uselessness with eligibility to remain, not to prove that a society generally practiced abandonment of the elderly.

3 | Volume One, Chapter 1 | Maine’s “From Status to Contract”

Henry Maine’s formulation arose in a particular age and from particular materials. It should not be treated as one path shared by all civilizations. This book borrows only the tension that status can sustain a place while also creating lock-in.

4 | Volume One, Chapter 1 | The statistical scope of the US Great Resignation

“Nearly 48 million” is the annual total of quit events in US labor statistics, not the same number of distinct individuals. Causes include changing jobs, wages, care pressures, working conditions, and other factors. The main text uses it only as a social signal that the occupational covenant is being reassessed.

5 | Volume One, Chapter 1 | Finland’s basic-income experiment

The experiment provided unconditional basic income to randomly selected unemployed people, principally observing employment and well-being outcomes. It was not designed to test the “sense of being needed.” Structurally asking why income security cannot automatically resolve social position does not amount to judging the experiment itself a failure.

6 | Volume One, Chapter 2 | East Asian care and mutual-aid practices

Children’s cafeterias, community meals, and time banks are not the same institution. Their origins, organization, scale, and policy backgrounds vary by region. They are parallel references for how care and relationships gain public visibility.

7 | Volume One, Chapter 2 | Boundaries in the use of Ubuntu

Ubuntu has multiple uses across languages, histories, and intellectual traditions. It cannot be summarized as one “African institution,” nor used to erase real tensions between individuals and communities. The main text takes up only the questions opened by its relational view of personhood.

8 | Volume One, Chapter 3 | Future computing architectures and governance boundaries

Quantum computing is a real research and engineering direction, but “quantum–AGI” is not an already-established unified technical object. The main text does not infer from superposition, entanglement, measurement, or particular demonstrations of quantum advantage that future AI must become unauditable or accelerate to a civilizational threshold. Nor does it assume a dedicated governance system. It retains only boundaries of stopping, takeover, audit, authorization, and responsibility that hold across technological forms.

9 | Volume One, Chapter 5 | The future memory scene in Nairobi in 2050

“Nairobi in 2050” is a future thought scenario, not a present institutional fact. It tests how memory can be carried forward while rights, qualifications, and governing authority cannot become automatically hereditary through ancestral records.

10 | Volume One, Chapter 5 | The nuclear-waste repository scene in 2055

This is a future thought scenario testing the relationships among intergenerational responsibility, long-term risk, and successors’ reassessment. It corresponds to no specific current project.

11 | Volume One, Chapter 6 | Epistemic boundaries of the Fermi question and the Great Filter

The Fermi question rests on the tension between models in which technological civilizations seem likely to leave detectable traces and the absence of confirmed, undisputed evidence of extraterrestrial technological civilizations. The Great Filter is a class of explanatory frameworks, not an established fact. Cosmic silence cannot prove widespread self-destruction among other technological civilizations or rank civilizations by maturity. The main text uses this open question only to remind us that long-term survival has no cosmic guarantee.

12 | Volume One, Chapter 6 | Petrov and a false nuclear early warning

On September 26, 1983, the Soviet satellite early-warning system produced a false alarm. Duty officer Stanislav Petrov did not treat its warning directly as a confirmed attack. The event is often simplified to “one man saved the world,” but the actual nuclear command chain, subsequent verification, and historical causation are more complex. The main text uses it only to show that high-consequence warnings need questioning, independent review, and room to stop. Machine signals cannot automatically equal orders to act.

13 | Volume One, Chapter 6 | Stress-testing the diffusion of high-consequence capabilities

Real barriers to different high-consequence technologies vary greatly, constrained by equipment, materials, expertise, access, experimental validation, defenses, and organization. This section offers only a directional stress test: If barriers to some capabilities keep falling, safety and responsibility structures must cover more subjects and nodes. It does not declare that “destructive capability has already become individual” or provide a uniform timetable.

14 | Volume One, Chapter 6 | Limits of the mitochondria and endosymbiosis analogy

The endosymbiotic origin of mitochondria is an important framework in modern evolutionary biology: An ancient bacterial lineage entered a host and became a eukaryotic organelle through long evolution. The main text uses this only to show that sustained dependence may change system capabilities. It derives no human rights, state sovereignty, political authorization, or obligation for society to “merge into one whole” from cellular evolution.

15 | Volume One, Chapter 6 | Indigenous peoples and local resource governance

Amazonian Indigenous peoples, Pacific Island societies, and pastoral nomadic societies are internally highly diverse. Resource relationships, prohibitions, seasonal governance, and ecological knowledge cannot be summarized as one “Indigenous institution of symbiosis.” The main text makes only a limited comparison: Many local traditions have brought future resource conditions into present decisions through relationships, seasons, harvesting restrictions, or intergenerational experience. It does not romanticize any group or replace modern legal rights and authority.

16 | Volume One, Chapter 6 | An extreme scenario of civilization-scale irreversible power

This is an extreme governance thought experiment. It does not discuss how destruction could be implemented. It tests only an institutional question: If misuse of a power could close everyone’s future, who may control it, and how can single-point overreach be prevented?

17 | Volume One, Chapter 6 | Earthrise and the overview effect

Apollo 8 astronaut William Anders took the famous Earthrise photograph in December 1968. The “overview effect” is a later concept describing the sense of wholeness and cognitive shifts some astronauts reported after viewing Earth from space. Such experiences can change a sense of scale. They do not prove that seeing Earth automatically produces a uniform political stance, moral conclusion, or global sovereignty.

18 | Volume One, Chapter 6 | The future scene of overoptimization

This is an extreme future thought scenario, not a prediction. It tests a structural risk: If a system persistently minimizes friction, waste, and uncertainty around a few goals, might subjects’ space for dissent, exit, experimentation, and redefining questions shrink with them?

19 | Volume Two, Chapter 7 | The “thief of time”

The “thief of time” is a literary metaphor for intertemporal externalization in which present actors realize gains early and leave risks or costs to successors. It is not a formal legal or economic term.

20 | Volume Two, Chapter 7 | The signal-to-noise analogy

The book borrows “signal-to-noise ratio” to illuminate relationships among information, evidence, and indicators. It does not calculate social value directly as a physical signal or establish a uniform social-value function.

21 | Volume Two, Chapter 8 | The 2018 Uber autonomous-driving test crash in Tempe

NTSB report HAR-19/03 records that the system first detected the pedestrian 5.6 seconds before impact but did not correctly predict the path. About 1.2 seconds before impact it determined a collision was imminent. Under the design then in use, collision-mitigation emergency braking was not executed in that situation; the system relied on the vehicle operator to intervene. The NTSB identified operator distraction as the probable cause, and contributing factors including deficiencies in Uber ATG’s safety risk management, operator oversight, and management of automation dependence, along with pedestrian behavior and insufficient Arizona oversight of testing. The book uses this to discuss distributed responsibility, not assign all responsibility to one actor.

22 | Volume Two, Chapter 8 | The 2023 Cruise pedestrian crash in San Francisco and disclosure

NHTSA’s 2024 consent order records that on October 2, 2023, a human-driven vehicle first struck a pedestrian into a driverless Cruise vehicle’s path. After the Cruise vehicle collided and initially stopped, it moved again, dragging the pedestrian roughly twenty feet. NHTSA found that Cruise’s one-day and ten-day reports failed fully to disclose the postcollision dragging. California’s DMV suspended Cruise’s driverless deployment and testing permits on October 24, 2023. The main text uses the case only to show that postincident disclosure belongs to the chain of responsibility.

23 | Volume Two, Chapter 8 | The Boeing 737 MAX and MCAS

The FAA's summary of its 737 MAX return-to-service review confirms that the original MCAS design could rely on one angle-of-attack sensor. With a persistently erroneous high-angle signal, MCAS could reactivate after pilots used electric trim; repeated commands could produce substantial stabilizer mistrim. Subsequent changes included using two angle-of-attack inputs and limiting repeated activation and MCAS authority. The book uses these confirmed system problems to discuss layered responsibility in high-consequence automation, not attribute the two crashes to a single software factor.

24 | Volume Two, Chapter 8 | The banality of evil and an analogy of technical responsibility

Since Hannah Arendt proposed the "banality of evil" in Eichmann in Jerusalem, Eichmann's motives, ideology, and Arendt's interpretation have remained subjects of scholarly debate. This book does not morally equate technical accidents with the Holocaust. It borrows only a limited question of responsibility: Can roles, procedures, and organizational divisions of labor weaken individuals' questioning of the consequences in which they participate?

25 | Volume Two, Chapter 8 | Independent financial audits as an institutional reference

The US Securities Act of 1933, Securities Exchange Act of 1934, and subsequent SEC rules established independent accountants' audits as important capital-market institutions. The main text borrows only the logic that interested parties cannot monopolize self-validation. It does not propose copying the accounting industry's structure for AI audits or appoint any particular auditor the natural gatekeeper of high-risk AI.

26 | Volume Three, Chapter 9 | Limits of the Costa Rica comparison

Historical assessments of Costa Rica after abolishing its standing army should use verifiable sources, defined indicators, and explicit periods. This book does not adopt unverified dramatic narratives such as "it was widely mocked at the time," nor treat superlatives such as "most stable, prosperous, and green" as facts needing no specified basis of comparison.

27 | Volume Three, Project Coordination | CERN

CERN is an international scientific organization established by member states through an international treaty. The main text uses it as a reference for sustained transnational scientific cooperation, not as a direct equivalent of this book's proposed project-coordination model.

28 | Volume Three, Project Coordination | Gavi

Gavi has its own defined organizational, legal, and financial structures. The text uses it only as a reference for multi-actor global-health cooperation. "Project coordination" and "alliance" are analytical language, not a renaming of Gavi's institutional identity.

29 | Volume Three, Knowledge Cooperation | The "species' self-understanding"

The "species' self-understanding" is a literary expression for changes in the scale of human self-understanding brought by knowledge projects such as the Human Genome Project. It is not a formal biological concept.

30 | Volume Two, Chapter 10 | The fabricated 2023 image of an explosion at the Pentagon

In May 2023, a fabricated image claiming an explosion near the Pentagon spread on social media and was subsequently refuted by US officials and information from the scene. Markets briefly fluctuated during the trading day. The book uses it only to show that inexpensive synthetic content can trigger real reactions before verification, not that "one fake image can control markets."

31 | Volume Two, Chapter 10 | The Hong Kong deepfake videoconference fraud

In 2024, Hong Kong police disclosed that a company employee transferred funds as instructed after attending a videoconference with several deepfaked “colleagues and executives,” causing large losses. The book uses the case only to show that familiar faces, voices, and meeting settings can no longer carry high-risk authorization alone.

32 | Volume Two, Chapter 10 | Biden-voice robocalls in New Hampshire

Before the 2024 New Hampshire US primary, robocalls imitating Joe Biden’s voice discouraged some voters from taking part. Federal and state enforcement, regulatory action, and penalties followed. The case shows only that generative voice forgery has entered political communication and impersonation.

33 | Volume Two, Chapter 10 | Capability boundaries of C2PA and Content Credentials

The C2PA specification aims to provide verifiable provenance and history for digital content, explicitly requiring that the specification itself make no “good or bad” value judgments about provenance assertions. It can verify the binding between assertions and assets, as well as signatures and tamper status. It cannot prove that real-world facts photographed, described, or interpreted are necessarily true.

34 | Volume Two, Chapter 10 | Paid Twitter/X blue checks and impersonation

After Twitter introduced a new paid blue-check scheme in 2022, several impersonation incidents occurred, including an account impersonating Eli Lilly and falsely stating that insulin was now free. This book does not attribute subsequent market-price movements to one post. It discusses only the trust risk when the public’s understanding of verification diverges from the conditions for obtaining a badge.

35 | Volume Two, Chapter 10 | Starling Lab and digital-evidence provenance chains

Starling Lab and partners have explored cryptographic signatures, hashes, distributed preservation, and camera-origin authentication to maintain provenance and chains of custody for journalism and potential legal evidence. Such technologies help verify whether materials changed and the chains of assertions they came from. They cannot alone establish the full facts beyond the frame or legal responsibility.

36 | Volume Two, Chapter 10 | NIST’s first post-quantum cryptography standards

In August 2024, the US National Institute of Standards and Technology approved and published FIPS 203, FIPS 204, and FIPS 205, its first three post-quantum cryptography standards, covering key encapsulation and two digital-signature schemes. The book discusses migration readiness on this basis, without presenting the quantum threat as a countdown to a definite year.

37 | Volume Two, Chapter 11 | Alaska’s Permanent Fund and resident dividends

Alaskan voters established the Permanent Fund by constitutional amendment in 1976. Following subsequent legislative and judicial adjustments, the resident-dividend system paid its first equal dividends in 1982. The text uses it only to show that institutionalized shared benefits can form when specific public-resource revenues, long-term funds, and beneficiary eligibility have clear legal grounds. It does not prove all assets should pay dividends to everyone.

38 | Volume Two, Chapter 11 | Norway’s Government Pension Fund Global

Norway’s Government Pension Fund Global manages state oil and gas revenues over the long term for present and future generations. It operates under Norway’s laws and governance structures and does not pay direct universal dividends to residents. It is used only as an institutional reference for carrying resource revenues across time and public long-term saving.

39 | Volume Two, Chapter 11 | Finland's basic-income experiment

Finland's statutory basic-income experiment of 2017–2018 randomly selected two thousand recipients of unemployment-related benefits to receive €560 unconditionally each month. Kela's final assessment found relatively small overall employment effects; participants reported several more positive outcomes in life satisfaction, psychological stress, financial well-being, perceived bureaucratic burden, and confidence in the future. The experiment does not prove basic income suits every country or answer all questions of work motivation and meaning.

40 | Volume Two, Chapter 11 | GiveDirectly's long-term basic-income research in Kenya

GiveDirectly began long-term basic-income research in Kenya in 2017. The long-term group was scheduled to receive monthly cash for up to twelve years, compared with shorter-term cash, lump-sum transfers, and controls. Interim research during the pandemic found cash buffered some hunger, health problems, and psychological stress, while external shocks still affected business and income. This cannot mean "security makes everyone work harder" or "human nature has been proved."

41 | Volume Three, Chapter 12 | The actual state of returns to education

"Returns to education have fallen to zero" is not generally the case today. Returns vary greatly by country, class, field, and educational level. The text discusses only how AI's changes to some knowledge tasks require recalibrating the relationships among educational content, occupational returns, and subjects' capabilities.

42 | Volume Three, Chapter 13 | The Fermi question and the Great Filter

Neither the Fermi question nor the Great Filter provides empirical proof for Civilization Leap's institutional conclusions. The first arises from the open question of why other technological civilizations remain unconfirmed; the second is a class of hypothetical frameworks. They serve only to place long-term survival risks on a broader timescale.

43 | Volume Three, Chapter 13 | Two kinds of failed future

The two futures in this section are thought scenarios extrapolating real trends to extremes, not predicted outcomes. They compare possible consequences of different capabilities, institutions, and competitive structures.

44 | Volume Three, Chapter 14 | The future supply-chain scene

This is a thought scenario for an institutional prototype, not a description of actual deployment in Shanghai or elsewhere. The demand-prediction example does not presuppose access to residents' health data. Narrative quantities serve neither actual prediction nor theoretical proof.

45 | Volume Three, Chapter 14 | A "civilization of flows"

"Civilization of flows" is a literary formulation, not a formal institutional concept or a claim that all resources should be centrally dispatched. Property rights, use rights, public authorization, market transactions, and privacy boundaries remain distinct.

46 | Volume Three, Chapter 15 | In Search of a Voice

The characters, institutions, interfaces, patents, matching, and subsequent events in In Search of a Voice belong to a future thought scenario, not a record of a real project. It tests boundaries among wishes, capability, AI matching, identity labels, and real relationships.

47 | Volume Three, Chapter 15 | “Exchanging a wish for admission”

“Submit a wish” is a literary device in *In Search of a Voice*, not a real public institution, universal trading rule, or requirement to disclose private wishes. Any similar real-world interaction should remain voluntary and protect privacy.

48 | Volume Three, Chapter 15 | Narrative numbers in *In Search of a Voice*

The numbers of people, wishes, and completed wishes, distances, and times in the story serve narrative functions. They are not actual platform data or evidence of the effectiveness of Dream Realization Park-type prototypes. Real trials must assess completion rates, failure rates, dependence, privacy, and negative consequences using real data.

49 | Volume Three, Chapter 15 | The “memory bank”

The “memory bank” is a future thought scenario discussing older people’s experience, digital memory, relationships carried forward, privacy, and subjects’ rights. It does not predict the institution will inevitably appear.

50 | Volume Four, Chapter 17 | The “hand and mind”

The “hand and mind” is only a literary metaphor for speed mismatch. Civilization has no unified brain. Actual judgment emerges from different subjects, institutions, professions, and chains of responsibility together.

51 | Volume Four, Chapter 18 | The Venetian Republic as a reference for institutional lock-in

The Venetian Republic’s long stability and eventual end involved republican governance, oligarchization, trade patterns, military and geopolitical pressures, and other factors. The book uses it only as an analytical lens on how long-successful structures can produce path lock-in, not a single-cause explanation of rise and fall.

52 | Volume Four, Chapter 19 | US security narratives after 9/11

After September 11, 2001, US security policy and public narratives changed profoundly, intertwined with the wars in Afghanistan and Iraq, detention, surveillance, domestic politics, and other factors. The text discusses only how crisis narratives affect risk priorities, policy justification, and exit from temporary measures. It does not attribute subsequent events to a single narrative.

53 | Volume Four, Chapter 19 | “Narrative gravity”

“Narrative gravity” is a literary metaphor, not a measurable physical variable of civilization or an empirical law that stronger shared narratives produce greater civilizational stability.

54 | Volume Four, Chapter 27 | Venetian glassmaking and Murano

Around 1291, Venice concentrated glassmaking on Murano for reasons including fire prevention, industrial management, and technological secrecy. Later prosperity and changing competition also reflected trade, technological diffusion, policy, and markets. The text uses this only to discuss how long-term protection can create path dependence, not a single-cause narrative that closure inevitably causes decline.

55 | Volume Four, Chapter 30 | The musician and instrument

This analogy describes only present functional cooperation, without permanently fixing future AI as a tool. If new evidence of subjecthood emerges, ethical and institutional positions must be reassessed.

56 | Volume Four, Chapter 30 | The “cognitive organ”

“Cognitive organ” is a limited metaphor for AI extending analysis, memory, and coordination. Civilization is not an organism with a single brain. The metaphor itself creates no data rights, executive authority, or sovereignty.

57 | Volume Four, Chapter 31 | The Fermi question and long-term civilization

The Fermi question and related Great Filter explanations remain an open space of competing hypotheses. Cosmic silence cannot prove technological civilizations tend to self-destruct or establish ranks of civilizational maturity. The main text uses it only as a thought-based stress test of the fragility of long-term survival.

58 | Volume Four, Chapter 32 | The image of “keeping watch”

“Keeping watch” is only an image of responsibility: Beyond others’ everyday lives, someone still remembers risks, boundaries, commitments, and long-term duties. It corresponds to no uniform historical class and constitutes no uniform personality education should form.

59 | Volume Four, Concluding Chapter | A “silent ending”

A “silent ending” is a literary classification of risk, contrasting visible destruction with less perceptible long-term lock-in. It does not claim these forms of risk are historically more common than war, disaster, or other causes.

60 | Volume Five, Preface | Distant forms of life and intelligence

Extreme longevity, digital life, personality uploading, reincarnation, multiple worlds, and fusion of human and machine subjects are not currently established unified facts. References in this volume are future stress tests or thought experiments. They cannot retroactively change present human rights or constitute mandatory routes of Civilization Leap.

61 | Volume Five, Chapter 3 | “Reincarnation”

“Reincarnation” here is only a thought experiment in resetting stages. It asserts no religious truth and proposes no deliberate manufacture by future civilization of amnesia, death, randomness, or irreversible costs.

62 | Volume Five, Chapter 4 | “Multiple worlds”

The volume treats “multiple worlds” not as a route of civilizational development, but as a comparative thought experiment concerning digital environments, identity migration, and differing rules. More worlds, greater complexity, and immersion do not automatically increase meaning.

Further Reading | From Understanding Civilization to Seeing How It Lives

Having read Civilization Leap: The Path of Civilization in the AGI Era (All Five Volumes), readers have completed the central civilizational argument of this body of thought: why a leap is necessary, what it must change, how civilization can operate over the long term, and why “Freedom from Survival Anxiety → Happiness in Life → Greater Meaning in Life” must remain a progression of goals. Continued existence remains a civilizational baseline that must never be lost; it is not recast as a fourth goal.

If the most natural next question is how such a civilization would actually live once it enters ordinary people’s work, families, relationships, cities, and collaboration with AI, turn first to the Living Civilization trilogy.

I. First, Toward Life: Enter the Living Civilization Trilogy

Living Civilization is not a “popular edition” of the five volumes, nor does it retell them in different words. The five volumes establish causation, goals, and structures at the civilizational level; Living Civilization brings those structures into real life, where they can be experienced, questioned, and compared.

Volume One, Principles for the Operation of a Future Civilization, focuses on how new civilizational relationships work. Volume Two, Real Life in a Future Civilization, places those relationships in ordinary people’s daily lives. Volume Three, The Formation of a New Civilization, goes on to answer another question: this life does not descend from the sky, so how does it gradually take shape through today’s institutions, technologies, human hearts, and real-world choices?

Moving from the five volumes into the trilogy therefore does not mean descending from “theory” to “stories.” It means proceeding from “why and what” to “what would happen if we actually lived within it.”

II. To Pursue the Questions of Today’s Crises: Return to What Is Happening in the AI Age

Continue with Humanity’s Civilizational Choice in the AGI Era and The Only Path. Through the crisis of white-collar work, agents, control, military AI, and the formation of rival blocs, the former asks what kinds of intelligent capability are amplifying the civilizational logic humanity actually follows today. The latter goes further into the structural contradictions among national security, deterrence, autonomy through severing connections, preemptive self-protection, and shared survival.

III. To Pursue Real-World Institutions and Engineering: From Civilizational Structure to Verifiable Implementation

Continue into the contribution-value system, resource pools, the separation of three rights, Spark, Nebula, and the related papers on experiments and cities. They investigate how real contributions can be given institutional support, how resources can be coordinated within boundaries of rights and authorization, how ordinary people’s capabilities can reach real needs with AI’s assistance, and how new institutional relationships can be tested within actual jurisdictions and the conditions of each stage.

For AI/AGI safety and high-consequence human judgment, continue into the Common Window Protocol, AGI-COS, and Human-COS. Their task is to protect the conditions for judgment, responsibility, audit, rollback, and renewed choice before irreversibility. They do not replace the civilizational goals of the five volumes.

IV. To Pursue Theoretical Foundations and Specialist Scrutiny

For deeper foundations, turn to T-10, Foundations for Interdisciplinary Dialogue and the Boundaries of Propositions; The Dynamics of Meaning; Thought Leap; and the papers related to Emergent Meaning Cosmology. For engagement with specialist communities, turn to the “1 + 16 + 2” academic Canonical system, CL-ACA-005–023. Deeper theory must not leap across levels to become empirical proof of real-world institutions. Academic peer scrutiny is likewise not a higher authority ruling over the civilizational Canonical. It is an interface through which facts, concepts, methods, and counterexamples feed back to recalibrate it.

V. If You Disagree

This theory should not use further reading to keep readers inside a closed system. New facts, counterexamples, failed trials, risks of power, specialist criticism, and better alternative explanations should all be able to return to the corresponding source manuscript, triggering local revision or even a renewed review of the Canonical. The most valuable participation is not agreement, but making errors visible earlier.

The five volumes answer why civilization must make a leap; Living Civilization answers how people genuinely live after that leap. Ultimately, these two paths of reading do not prove each other. They test each other.

Unified public entry point: www.civitas.top

(End of the Complete Work)

A Further-Reading Guide for Your Next Questions

This is not a list of “related reading.” It is a way to pursue the questions that follow from this work.

You have completed the central civilizational argument. The next step is not to read the theory again, but to put it to the test in real life, in today’s AI realities, and in the correction of real-world errors.

Toward Life

Living Civilization: The Complete Trilogy (Public Collected Edition)

Move from civilizational theory to seeing how it takes shape in real life.

[Living Civilization: The Complete Trilogy \(Public Collected Edition\)](#)

Toward Present Realities

Humanity’s Civilizational Choice in the AGI Era: What Are We Teaching Future Intelligence?

Return to today’s AI realities and civilizational choices.

[Humanity’s Civilizational Choice in the AGI Era: What Are We Teaching Future Intelligence?](#)

Toward Error Correction

Seeing Earlier What Would Otherwise Be Seen Only Later: A Public Introduction to Human-COS

Enter real-world judgment and error correction.

[Seeing Earlier What Would Otherwise Be Seen Only Later: A Public Introduction to Human-COS](#)

An Entry Point for Questions, Counterexamples, and Revision

New facts, counterexamples, failed tests, new capabilities, new attack surfaces, or changes in the conditions of a stage should all enter renewed review within the relevant family of papers. Further reading must not be used to preserve existing conclusions.

[Complete Library](#)

[Main Public Entry Point](#)

Main text unchanged | Library navigation page added only | Civilization Leap Research Group | civitas.top