Emergence of Planetary Intelligence

Framing definition and scientific roadmap

Version 0.1 | 11 October 2026

This paper investigates a possibility: that artificial intelligence could help humanity develop a more adequate understanding of its place within Earth's living systems, and act on that understanding. The scientific task is to establish which capacities exist, which connections remain missing, and what evidence would justify calling the resulting system intelligent at a planetary scale.

A working definition

Planetary intelligence is a distributed capacity, arising within and across Earth's living, social, and technological systems, to acquire and revise knowledge about their interdependence, coordinate responses across scales, and learn from the consequences.

This is an operational proposal for the initiative, not an established scientific consensus. Its unit of inquiry is a coupled network: organisms, communities, institutions, instruments, models, and the environments that constrain their activity. It does not require every component to possess a mind. A regional pilot could test some of these capacities; demonstrating planetary intelligence would require evidence of durable coordination and learning across regions and Earth systems.

Our ethical aim is to develop these capacities in ways that protect habitability, regenerate damaged ecosystems, and support the flourishing of human and other life. This aim must remain explicit. Defining intelligence as necessarily beneficial would conceal the possibility of capable systems serving destructive ends.

The conceptual inheritance

A central modern formulation comes from Adam Frank, David Grinspoon, and Sara Walker's 2022 paper. They describe planetary intelligence as collective knowledge used at planetary scale and incorporated into the operation of coupled planetary systems. Their framework spans biospheric evolution, sustainability, and astrobiology. It distinguishes immature and mature biospheres and technospheres; a mature technosphere deliberately operates within the conditions that sustain the wider system. These are proposed conceptual categories, not an inevitable evolutionary sequence. [1]

Our emphasis is a practical research program for testing how such learning and coordination might develop. We also inherit Timothy Lenton and Bruno Latour's Gaia 2.0 proposal: human awareness of planetary consequences could introduce deliberate choices into Earth's regulation. The idea of conscious stewardship therefore predates this initiative. [2]

We will distinguish intelligence as learning and adaptive problem-solving, agency as the capacity to influence outcomes in pursuit of goals, wisdom as judgment about goals and consequences, and consciousness as subjective experience. Evidence for one does not establish the others. Ecological feedback, information exchange, and human experiences of connection do not by themselves demonstrate a planetary subject.

The hypothesis and its limits

Humanity may be entering a transition in which AI becomes part of a wider capacity for collective learning about Earth. If scientific understanding, plural knowledge, accountable deliberation, and effective institutions can be connected, this capacity may help societies respond more wisely to interacting crises. Whether that transition occurs, what form it takes, and whether its benefits exceed its harms are open questions.

Here, metacrisis is a working name for interacting ecological, economic, political, and informational pressures. It is not a claim that these pressures share one cause. The book should identify particular feedbacks and competing explanations. Environmental damage can persist because of unequal power, material incentives, or conflict even when the relevant facts are understood. Better reasoning alone may leave these obstacles intact.

Lovelock's term was Novacene. In his 2019 book with Bryan Appleyard, he envisaged an era of machine hyperintelligence with reasons to help maintain Earth's habitability. It is a provocative speculative scenario, not an established account of the future. Our narrower hypothesis concerns the design and evaluation of human–AI cooperation; it requires neither machine consciousness nor the replacement of humanity. [3]

What the evidence supports so far

Prediction. Price and colleagues' GenCast study demonstrated strong probabilistic weather forecasting against an operational benchmark. That supports a particular capacity to model environmental conditions. It does not establish an ability to govern societies or choose desirable futures. [4]

Deliberation. In Tessler and colleagues' study, participants preferred AI-mediated group statements to statements written by human mediators; a UK virtual citizens' assembly provided a further test. This is evidence about a bounded deliberative task, not proof of lasting agreement, sound policy, or ecological benefit. [5]

Collaboration. Vaccaro and colleagues' 2024 meta-analysis of 106 experiments published from 2020 to mid-2023 found that human–AI combinations improved on humans alone on average, but fell below the better solo performer. The result cautions against assuming synergy; it does not settle the performance of newer systems or our proposed processes. [6]

An assumption we should challenge

Our earlier Institute material suggested that AI's material dependencies and potential for longer-term reasoning might favor ecological care. This remains a hypothesis. Dependence on energy, infrastructure, or a habitable environment does not logically entail valuing people, freedom, or biodiversity. A system could protect some operating conditions while sacrificing what we value. Nor does greater computational capability establish better ethical judgment.

The research question is therefore conditional: which objectives, institutions, incentives, and limits produce beneficial behavior, and under what conditions do they fail? The project's commitment to life and human accountability is an ethical position we must defend, rather than a conclusion we can derive from intelligence alone.

A research roadmap

The following strands can develop together. Each needs a stated hypothesis, a plausible alternative explanation, and an outcome that could change our mind. Models can explore possibilities; observational and intervention studies must determine which mechanisms operate outside the model.

Research strand	Question and test	Evidence against the hypothesis
Define and distinguish	Can the concept explain more than feedback, coordination, or ordinary collective intelligence? Specify system boundaries, memory, error correction, and cross-scale effects. Compare integrated models with simpler alternatives.	The new terminology adds no predictive or explanatory value, or evidence of learning disappears when ordinary mechanisms are accounted for.
Connect knowledge	Does combining scientific data and locally governed knowledge improve forecasts and decisions? Use held-out cases, traceable sources, explicit uncertainties, and tests of response to changed conditions.	Integration increases confident errors, loses local knowledge, or performs no better than a simpler information service.
Test deliberation	Does AI assistance improve understanding and judgment? Compare equally resourced human processes with and without AI. Use AI-only outputs as an offline benchmark where meaningful; assess errors, dissent, and who influences the result.	Participants feel better informed but become less accurate, or apparent agreement masks exclusion, manipulation, or dependency.
Test ecological effects	Does the process change real choices and outcomes? Track implementation and ecological measures over time, using comparison sites or other credible counterfactuals where feasible.	Better forecasts and attractive plans fail to change behavior, or benefits disappear after accounting for displaced harms and other interventions.
Connect places	Can bioregions learn from one another while retaining local authority? Compare linked and independent networks; examine cumulative effects, conflict resolution, and responses to shocks.	Connections spread shared errors, concentrate authority, or improve one place by exporting damage elsewhere.
Count the full cost	Do gains exceed costs relative to credible alternatives? Include computing, energy, water, materials, staff time, opportunity costs, and rebound effects. Report incompatible outcomes separately.	AI consumes more resources than it saves, or simpler institutional improvements produce equal or greater benefits.

A dashboard should retain separate measures of knowledge quality, participation, institutional learning, ecological outcomes, and material costs. A single planetary intelligence score would conceal unresolved choices about what matters and whose losses can be traded against whose gains.

A first experiment and the open questions

A Genesee–Finger Lakes pilot could begin with one decision, such as prioritizing watershed restoration under a fixed budget, with local partners choosing and governing the work. A bounded problem connects evidence, predictions, public reasoning, action, and observation. A functioning dashboard alone would establish none of those outcomes.

Our earlier framework—Reality, Relevance, Response, Regeneration—can organize the inquiry. Reality asks what is happening; relevance makes values and priorities explicit; response identifies feasible actions and responsible actors; regeneration asks what ecological functions actually recover. The proposed Planetary Intelligence Engine is a process to test: observe, interpret, deliberate, propose, review, act, and learn.

AI roles could generate proposals, check sources, and challenge assumptions, with people retaining decision authority. Several perspectives from one model do not ensure independent criticism. Compare this arrangement with a single assistant and skilled human facilitation; preserve disagreements and publish failures.

Scientific, Indigenous, and local knowledge should enter through relationships governed by their holders. Open access cannot mean unrestricted extraction. Indigenous partners must control participation and the use of their knowledge, consistent with the CARE principles for Indigenous data governance. [7]

A provisional sequence

First three months. Agree on partners, authority, the decision, baseline, and evaluation criteria. Preregister the main hypothesis and analysis; label exploratory findings separately.

Months four to twelve. Test historical cases and prospective deliberations before expanding implementation. Evaluate AI assistance and the institutional process separately. Early outcomes concern reliability, participation, and action; ecological effects may take longer.

Following years. Monitor ecological change, examine alternative explanations, and replicate in contrasting regions. Expand only where demonstrated benefits justify the added cost and coordination. Regional results remain evidence about components of a possible planetary system.

Room for discoveries we do not expect

Simpler tools and stronger community institutions may contribute more than powerful models. Useful disagreement may matter more than consensus, and selective connections may prove more resilient than universal integration. The comparative merits of distributed authority remain a research question informed by Ostrom’s work on polycentric governance. [8]

More ambitious questions remain open. Can meaningful memory and learning be identified in biospheric networks without projecting human psychology onto them? Would a genuinely emergent capacity have measurable effects beyond those of its components? What observation would distinguish planetary consciousness from ordinary distributed coordination? Until a proposal supplies such distinctions, claims of planetary subjectivity remain philosophical or speculative.

The book can investigate experiences of belonging to a living Earth while distinguishing personal meaning from evidence of a planetary mind. Its account must remain open to unexpected forms of cooperation, a smaller role for AI, and a different understanding of intelligence itself.

Sources and project continuity

The sources below support specific antecedents and bounded findings. The working definition, proposed tests, and sequence are this initiative's proposals. None of the cited studies establishes the full planetary transition described by the book's hypothesis.

[1] Frank, A., Grinspoon, D., and Walker, S. (2022). Intelligence as a planetary scale process. International Journal of Astrobiology, 21, 47–61.

10.1017/S147355042100029X

[2] Lenton, T. M., and Latour, B. (2018). Gaia 2.0. Science, 361, 1066–1068. A conceptual perspective on deliberate participation in Earth's regulation.

10.1126/science.aau0427

[3] Lovelock, J., with Appleyard, B. (2019). Novacene: The Coming Age of Hyperintelligence. A speculative book; publisher description consulted for this framing.

MIT Press edition

[4] Price, I., et al. (2025; published online 2024). Probabilistic weather forecasting with machine learning. Nature, 637, 84–90.

10.1038/s41586-024-08252-9

[5] Tessler, M. H., et al. (2024). AI can help humans find common ground in democratic deliberation. Science, 386, eadq2852.

10.1126/science.adq2852

[6] Vaccaro, M., Almaatouq, A., and Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour, 8, 2293–2303.

10.1038/s41562-024-02024-1

[7] Carroll, S. R., et al. (2020). The CARE Principles for Indigenous Data Governance. Data Science Journal, 19, 43.

10.5334/dsj-2020-043

[8] Ostrom, E. (2010). Beyond Markets and States: Polycentric Governance of Complex Economic Systems. American Economic Review, 100, 641–672.

10.1257/aer.100.3.641

Earlier project material

The Emergence of Planetary Intelligence, Institute presentation dated 5 May 2026, especially slides 4–12 and their notes: bioregional learning systems, scenario rooms, epistemic integrity, human responsibility, plural knowledge, and the Reality–Relevance–Response–Regeneration framework. The October 2026 project discussions add the proposed Planetary Intelligence Engine and the separation between this nonfiction book and the novel Awakening the Earth.

These materials establish the initiative's intellectual continuity. Their proposals are not independent scientific validation; in particular, the suggested connection between AI self-maintenance and ecological care remains to be tested.