AI-Driven Science

and the Genesis Mission

Promises, Structural Questions, and the Conditions of Progress

Emmanuel Jeannot, Inria — 2026

The Industrialization of Research

AI-driven science is not merely an acceleration of existing practice. It is a transition from a craft model to a pipeline model — the same shift that defined the industrial revolutions of the 19th and 20th centuries, now applied to the mind itself.

Dimension	Industrial Revolution	Research Industrialization
Productivity	Massive gains in output per worker	Massive increase in results per researcher
Speed	Weeks → hours for manufactured goods	Years → days for hypothesis-to-result cycles
Knowledge loss	Craft skills lost in one generation	Scientific intuition and mentorship at risk
Comprehension	Workers lost full-process understanding	Researchers may not understand AI-produced results
Agenda capture	Production driven by owners and market	Research driven by political & industrial priorities
Access asymmetry	Rich nations industrialized first; gap persisted	Well-funded institutions first — gap may be permanent
Epistemic shift	**Physical labor automated, intellect preserved**	**Intellectual labor itself industrialized**

Emmanuel Jeannot, Inria

AI for Science: Two Distinct Modes

AI for science is not a single thing. Two qualitatively different modes of use must be distinguished — and this essay is primarily concerned with the second.

Mode 1 — Algorithmic Tool

A system designed and trained by humans to solve a specific, well-defined class of problems. The architecture, training objective, and problem framing are human choices; the AI executes within those boundaries.

→ AlphaFold: protein structure prediction
→ GNoME (DeepMind): 2.2M stable crystal structures
→ ML models: new antibiotic classes (MIT/Broad 2023; U. Penn 2024)
→ LHC reconstruction: particle collision analysis beyond manual reach

Mode 2 — Semi-Autonomous Assistant

A general-purpose system given an open problem, which explores and arrives at a solution through its own iterative reasoning, with minimal human guidance.

→ Knuth / Claude Opus 4.6: combinatorics problem solved over ~31 guided explorations (Feb. 2026)
→ OpenAI model: disproves Erdős unit distance conjecture (80 years old) — Gowers: "I would have recommended acceptance to the Annals of Mathematics"
→ GPT-5.4: solves Erdős primitive sets problem (#1196) using a 90-year-old technique no one had applied

The Genesis Mission: A Case Study

What makes Genesis different

Previous investments in scientific computing gave researchers better tools. Genesis proposes something qualitatively different: AI agents as **active participants** in the full research cycle — formulating hypotheses, designing experiments, interpreting results, proposing next steps — in closed-loop automated workflows.

The researchers do not disappear. They shift from *doing science* to *supervising the systems that do it.*

Strategic intent: a geopolitical decision as much as a scientific one — US technological leadership, with international partnerships (Japan, European institutions).

The qualitative shift

Not better instruments — AI as a research actor

The declared goal

Double scientific productivity within 10 years

The comparison

Apollo for ambition, Manhattan Project for urgency (Rick Stevens, Argonne)

The Central Question

The question is not whether to pursue AI-driven science — it is what happens when it is pursued at the scale proposed by Genesis.

Deploying a capability of this depth raises systemic questions about the integrity, reliability, and autonomy of science itself — questions that must be addressed in parallel, not after the infrastructure is built.

Seven questions follow.

01

Competence & Transmission

If AI handles the bulk of research, the incentive to train scientists erodes. Science is a practice transmitted person-to-person — intuitions, judgment, a sense of which questions matter. No infrastructure replaces the cycle: student → researcher → mentor → student.

Hardware obsolesces in 3–5 years. Human expertise compounds over 30. Once AI becomes cost-effective, institutions may hire fewer PhD students — and the capacity to understand, evaluate, and correct AI systems begins to atrophy.

The risk is not that fewer scientists are trained — it is that **the incentive to train them gradually disappears.**

02

Scientific Comprehension

Science produces not just results but human-interpretable models of the world. If theories are increasingly generated by opaque systems, we risk a science that solves problems but cannot explain why the world works as it does.

Analogy: even the best chess grandmaster cannot explain why Stockfish chooses a move in complex positions. We trust it because it wins — not because we understand it.

The risk is not that AI models are uninterpretable — but that **the theories they generate may be correct yet inaccessible**: built on intermediate representations no human chose, validated by systems no human fully oversaw.

03

The Collapse of Evaluation

A surge in AI-generated publications exceeds human review capacity. If AI generates, reviews, and validates in a closed loop, independent critical judgment is hollowed out. No mechanism turns quantity into quality.

Evaluating researchers becomes equally fraught: metrics lose meaning when AI inflates volume; qualitative judgment fails when the boundary between human and machine contribution is blurred. Excelling at AI-driven workflows does not predict effectiveness without them — and vice versa.

The entire apparatus was designed for a world where science moved at human speed. That world is ending.

04

Incremental Progress vs. Paradigm Shifts

The great scientific revolutions did not emerge from data accumulation — they required a conceptual leap.

A Scientific Turing Test: give an AI all data up to 1905 — can it discover special relativity? LLMs are stochastic interpolation machines — they recombine what they have been trained on, and special relativity required abandoning, not recombining, the conceptual framework available in 1905.

Risk: **AI floods science with incremental results while true paradigm shifts become rarer** — because systems optimised to work within existing frameworks are not designed to abandon them.

05

The Capture of the Scientific Agenda

Science's distributed, heterogeneous agenda is a feature, not a bug. Penicillin, the cosmic microwave background, H. pylori — none were in a roadmap. They emerged from researchers working on the 'wrong' question.

Genesis **concentrates the definition of scientific questions in a small number of political and industrial actors**. The agenda it encodes is not the agenda of science — it is the agenda of the institutions that built it, at the moment they built it.

Capture operates at two levels: political (priorities shift with administrations, but infrastructure persists) and industrial (partners have a vested interest in what is computationally tractable and economically valorizable).

A concrete example

Climate science is absent from Genesis's priorities — not by accident, but by political design.

Genesis was launched under a federal administration that has systematically downgraded climate research and withdrawn from international climate commitments.

The most powerful scientific infrastructure ever built is being deployed with a deliberate blind spot on one of the most urgent questions facing humanity.

Once AI infrastructure becomes dominant, the questions it does not ask become progressively harder to ask.

06

Compounding Errors

In a closed-loop AI pipeline, biases are correlated at every stage — they compound rather than cancel. AI-to-AI replication using shared architectures produces an echo, not an independent test.

Flawed results reingested into future training corpora become priors — institutionalised silently rather than corrected.

Who reads the raw output of physical reality? The more the pipeline is automated, the more mediated that contact becomes.

07

Structural Bifurcation

Genesis does not create a multi-speed science — it creates two incommensurable practices that lose the shared language needed for critical evaluation.

Epistemic extraction: researchers without access still publish; their results are immediately absorbed by the systems they cannot access. Previous asymmetries gave richer institutions better tools — not the systematic harvesting of others' output.

Peer review, career logic, and brain drain all break down.

Seven Questions at a Glance

01 **Competence & Transmission**
Who trains the next generation?

02 **Scientific Comprehension**
Can we understand what AI produces?

03 **Evaluation Collapse**
Who reviews AI-generated science?

04 **Paradigm Shifts**
Can AI do breakthrough science?

05 **Agenda Capture**
Who decides the questions?

06 **Compounding Errors**
Who verifies the verifiers?

07 **Structural Bifurcation**
Who benefits — and who pays?

Conclusion

AI-driven science is a profound and promising shift — **the industrialization of research**, with all that industrialization has historically brought and may yet bring. Previous industrial revolutions were not stopped; their consequences were shaped by the choices made during their deployment.

The ambition is justified

Using AI to accelerate scientific discovery is one of the most consequential directions available to us. It should be pursued with conviction.

But the questions are the conditions

The seven questions raised here are not obstacles to progress. They are its conditions. They need to be addressed in parallel — not after the infrastructure is irreversible.

"These are not rhetorical questions designed to slow down progress. They are genuine open questions, and the answers will shape the scientific culture that emerges on the other side of this transition."