



AI in Science and Questions to Ponder
July 2026
ISSUED BY HETERODOX LABS

Scientia

Automated science is all the rage presently, as every month brings another headline proclaiming that artificial intelligence has made a scientific discovery or predicted some promising drug candidate or identified a new protein. Following such announcements comes the familiar refrain that science itself will soon be “solved,” as though it were a Rubik's Cube awaiting sufficient compute.

Given my tone, you can probably imagine how such headlines sit with me. Yet beyond abstract philosophical concerns, I think there are widespread and fundamental misconceptions that plague discourse around AI and the scientific enterprise.

One could argue that the project of creating the machine of discovery was initiated not in Silicon Valley, but nearly four centuries ago. In his *Discourse on the Method*, René Descartes envisioned science as a means to “render ourselves the masters and possessors of nature.” The Scientific Revolution and Enlightenment brought with them a growing confidence that science, and its practitioners, possessed privileged access to truth, and that through sufficiently rigid methodological standardization, the Scientific Method™, we might steadily unlock the universe until little mystery remained. Such ambitions were, in many respects, extraordinary successes. But this success breeds arrogance, and can lead one away from the holy path. And indeed, the reductionist impulses of the Enlightenment form much of the foundation for our present and rather narrow conception of scientific discovery.

Today, this inheritance has simply taken a new form. There exists a growing notion that with sufficient compute, larger language models, and enough data, one can uncover all of the universe's secrets. But what does it even mean to “solve science”? The question is, in many ways, preposterous. It is akin to asking whether we might someday solve art, solve religion, or solve love. Such notions do not really mean anything at all. Our interpretations of nature's patterns are not, as many in recent years have claimed, “settled,” nor should we expect them ever to be.

In many ways, it is through messy human forces like rhetoric and argument that scientific understanding has advanced—far more than many would like to believe. Science is not merely the accumulation of facts, but it is our primitive instinct to continually reinterpret them. In this sense, it bears some resemblance to a court of law. In certain cases, overwhelming evidence leaves little room for doubt. There is clear footage, abundant testimony, and a preponderance of evidence. In others, the matter remains uncertain, contested, and perhaps destined to remain unresolved. In such cases, humility is of greater scientific virtue than certainty.

Many sweeping proclamations are being made at this moment, particularly within the AI-for-science conversation, when the only truly obvious reality is that the future of science will remain what it has always been: messy, chaotic, deeply interpretive, and perpetually unfinished.

With that, how should we proceed? How ought we to carry ourselves epistemically? What assumptions should we be cautious of? And perhaps most importantly, what principles should we choose to imbue within the intelligent tools we are now building, and thus the kind of science they will ultimately encourage?

I raise these questions not to answer them, but in the hope that you will ponder them.

INSIDE HDX

AI + Discovery in Washington, D.C. – Last week we attended the AI + Discovery event in Washington, D.C., where we heard from folks across science, government, and industry, including representatives from the Department of Energy, National Science Foundation, Los Alamos National Laboratory, Oak Ridge National Laboratory, Anthropic, Foresight Institute, and others. Unmistakably, there is rapidly growing interest in AI for Science at the national level, a trend reflected in recent initiatives and reports such as Director Michael Kratsios' *Science: A Golden Age*. Throughout the week, we continued meeting with researchers, founders, and organizations working across AI for science, lab infrastructure, and scientific reproducibility, helping shape future collaborations and pilot opportunities.

Experiment Engine in an Applied Use Case – While Experiment Engine was designed with scientific laboratories in mind, we believe its underlying approach to procedural observability may extend beyond traditional research. Food, beverage, and pharmaceutical R&D all involve complex processes where transparency, documentation, and reproducibility are increasingly valuable.

To explore one such application, we've been documenting a brewing project in the lab using both Experiment Engine and conventional video. The goal is to explore how AI-assisted procedural documentation might create a new level of transparency for both producers and consumers. More on that, and perhaps a taste, at our annual event.

SAVE THE DATE

Our annual Heterodox event will take place August 29th in downtown Norfolk, Virginia. Join us for an evening of talks, demonstrations, drinks, food, and conversation as we share what we've been building over the past year.