
Research in 2035: The Industrialization of Discovery

Odysseus Kaloidas
UC Berkeley

Abstract

By 2035, research could depend more on directing AI agents across disciplines and less on personally mastering every specialized procedure. Agents are systems that carry out sequences of tasks toward a goal. If they become reliable enough, reusable processes could turn selected questions into sustained programs of analysis and experiment: an industrialization of research. Expertise would partly appear in the questions, examples, and standards used to guide that work. This is a scenario, not a dated prediction. It asks how researchers would learn, specialize, and develop judgment when one person's way of thinking could organize far more work than that person could perform alone.

1 Following a Question Across Fields

Why do two patients with apparently similar disease respond differently to treatment? Following that question may require molecular biology, statistics, clinical knowledge, and software engineering. A researcher can spend years acquiring another technique or finding the right collaborator. Imagine instead commissioning AI agents to retrieve evidence, implement analyses, compare explanations, and return unresolved choices. These systems carry out sequences of tasks toward a goal. In this 2035 scenario, the researcher follows the question across disciplines by directing work within them.

A person could follow a question further without first mastering every procedure needed along the way. Jones's account of the burden of knowledge connects accumulating knowledge with specialization and teamwork [4]. Agents could loosen part of that constraint by making specialized procedures easier to access and combine. They would also create a new educational emphasis: learning how to commission, configure, and understand work done through AI.

2 Borrowing a Method, and the Means to Use It

A borrowed method could arrive with a system capable of applying it. Paper2Agent, published in September 2026, converts papers and their associated software into callable tools for agents; its case studies include genomic and single-cell analyses [8]. A researcher can then use the method through conversation instead of first reconstructing how to run its software.

Virtual Lab offers a complementary example. It organizes agents with different scientific roles under an agent principal investigator, with human guidance, and reports computational design followed by experimental validation of protein binders [10]. This is a particular biomedical demonstration. It suggests how a researcher might assemble the needed specialties around a question instead of limiting the question to the available specialty. In a February 2026 interview, James Zou describes exploring different arrangements of agent discussions in parallel [11]. Even the organization of collaboration becomes something that can be varied computationally.

For an individual, interdisciplinarity could therefore mean more than reading broadly. It could mean conducting a study that joins several technical traditions. Giving instructions in a shared language could ease coordination. The researcher would still need to understand where the methods make

different assumptions. A predictive model, a causal explanation, and a clinically useful intervention remain different achievements, however smoothly agents connect their outputs.

3 Learning to Direct, Specializing Through Taste

Microsoft’s account of the emerging “agent boss” already describes workers building, managing, and delegating to agents [9]. Applied to research, the deeper shift concerns what people study. Training may devote more attention to specifying investigations, selecting methods, interpreting failures, and configuring capable systems, with less time required to master every implementation personally.

Two researchers could use the same underlying models and build quite different research programs. One might pursue a troublesome anomaly; another might look for a simplification that connects familiar results. Their choices would shape the literature, examples, criticism, and standards supplied to their agents. This is part of scientific taste: judging which questions, explanations, and standards are worth pursuing. Some of it can be expressed in instructions; some is hard to articulate and emerges through repeated correction.

Agents may also help formulate questions and revise the standards themselves; scientific taste need not remain an exclusively human contribution. A researcher’s expertise could meanwhile reach further through the systems they configure. But directing that work still requires subject knowledge. Bainbridge’s analysis of industrial automation identifies the difficulty of leaving people responsible for exceptional situations after automating ordinary operation [2]. In research, learning only to request an analysis could leave someone unable to recognize when its premises fail.

4 From Individual Projects to Discovery Production

A research group might build a process once and use it to investigate many related questions. Problems enter; agents search, implement, test, and revise; selected results proceed to further investigation. Compute becomes a way to purchase additional attempts, while improvements to the workflow can benefit subsequent projects. The factory analogy concerns organized, repeatable production of investigations. Its products need not be identical, and its organization can remain adaptive.

The idea of organized discovery predates today’s agents. Cockburn, Henderson, and Stern examine AI as a general-purpose method of invention that could reorganize research [3]. Kitano’s Nobel Turing Challenge imagines an engine that generates and verifies hypotheses at scale [5]. The AI Scientist demonstrates a bounded version within AI research, connecting idea generation, code, experiments, analysis, and manuscript production [6]. These precedents make the factory analogy concrete, while leaving open how far the process can extend beyond the problems it can currently formulate and check.

Suppose agents can execute thirty investigations a month but the available facilities can validate only ten. If each investigation needs validation, faster execution alone cannot sustain more than ten validated completions a month. That comparison holds the class of work and validation capacity fixed; redesigning experiments or expanding facilities can change it. Completion also differs from discovery: an impeccably conducted investigation can reject its hypothesis. Scientific taste affects which investigations enter the process and what their results mean, as well as how much work gets done.

Amodei’s optimistic *Machines of Loving Grace* likewise discusses limits from experimental duration, unavailable evidence, and the physical world [1]. Some constraints can themselves become targets of invention. Research industrialization could therefore improve both the work and the machinery that performs it, while remaining uneven across domains.

5 What the Factory Makes Room For

A researcher’s work could become broader while science as a whole becomes narrower. Agents could let more people work across fields while directing them toward the same accessible datasets and easily evaluated problems. Messeri and Crockett warn that AI can encourage scientific monocultures, a narrowing around shared methods and questions, and an exaggerated sense of understanding [7]. A factory optimized for measurable output could intensify that tendency. A system capable of cheap

exploration could also make neglected combinations affordable. Neither outcome follows from automation alone.

A researcher in this scenario would need to understand enough of each method to judge how it fits the question, without personally performing every step. Education would have to prepare people for that combination of breadth and dependence. Specialization could then appear as much in the investigations a person directs as in the procedures they can execute. One person’s way of thinking might organize more work than their hands and training could sustain alone. The unresolved question is how to develop that judgment when so much of the work through which it was once learned is done by agents.

References

- [1] Dario Amodei. Machines of loving grace. Author essay, October, 2024. URL <https://darioamodei.com/essay/machines-of-loving-grace>.
- [2] Lisanne Bainbridge. Ironies of automation. *Automatica*, 19(6):775–779, 1983. doi: 10.1016/0005-1098(83)90046-8. URL [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8).
- [3] Iain M. Cockburn, Rebecca Henderson, and Scott Stern. The impact of artificial intelligence on innovation. Technical Report 24449, National Bureau of Economic Research, 2018. URL <https://www.nber.org/papers/w24449>.
- [4] Benjamin F. Jones. The burden of knowledge and the “death of the renaissance man”: Is innovation getting harder? *The Review of Economic Studies*, 76(1):283–317, 2009. doi: 10.1111/j.1467-937X.2008.00531.x. URL <https://doi.org/10.1111/j.1467-937X.2008.00531.x>.
- [5] Hiroaki Kitano. Nobel Turing Challenge: creating the engine for scientific discovery. *npj Systems Biology and Applications*, 7:29, 2021. doi: 10.1038/s41540-021-00189-3. URL <https://www.nature.com/articles/s41540-021-00189-3>.
- [6] Chris Lu, Cong Lu, Robert Tjarko Lange, Yutaro Yamada, Shengran Hu, Jakob Foerster, David Ha, and Jeff Clune. Towards end-to-end automation of AI research. *Nature*, 651:914–919, 2026. doi: 10.1038/s41586-026-10265-5. URL <https://www.nature.com/articles/s41586-026-10265-5>.
- [7] Lisa Messeri and Molly J. Crockett. Artificial intelligence and illusions of understanding in scientific research. *Nature*, 627:49–58, 2024. doi: 10.1038/s41586-024-07146-0. URL <https://www.nature.com/articles/s41586-024-07146-0>.
- [8] Jiacheng Miao, Joe R. Davis, Yaohui Zhang, Jonathan K. Pritchard, and James Zou. Reimagining research papers as interactive and reliable AI agents. *Nature*, 2026. doi: 10.1038/s41586-026-11044-y. URL <https://www.nature.com/articles/s41586-026-11044-y>.
- [9] Microsoft WorkLab. How to be an agent boss. WorkLab, May 13, 2025. URL <https://www.microsoft.com/en-us/worklab/how-to-be-an-agent-boss>.
- [10] Kyle Swanson, Wesley Wu, Nash L. Bulaong, John E. Pak, and James Zou. The virtual lab of AI agents designs new SARS-CoV-2 nanobodies. *Nature*, 646:716–723, 2025. doi: 10.1038/s41586-025-09442-9. URL <https://www.nature.com/articles/s41586-025-09442-9>.
- [11] James Zou, Nathan Labenz, et al. Approaching the AI event horizon? part 1. The Cognitive Revolution, video interview, February 13, 2026. URL <https://www.youtube.com/watch?v=1V11KiNjXgg>. Discussion with James Zou, approximately 6:09–7:24.