THE AI NOTEBOOK · WHAT MATTERS AND WHY

Where AI breaks the pattern

AI Through the Lens of History · Part 4 of 5

By Dr Peter Fields

THE 2-MINUTE SUMMARY

Four features of AI make the historical comparisons less secure than they first appear.

Generative AI was tried quickly and can take on some of the junior tasks through which people learn a profession. Some systems can act on their outputs. Their mistakes can look plausible. The resources needed for the most expensive systems are concentrated in a small number of firms.

Human review can help when the reviewer has time, evidence and authority to intervene. Treating sign-off as sufficient is a weak safeguard.

Full article · About 4 minutes to read

Speed, and the first rung

AT&T installed its first dial exchange in 1919. Its last manual exchange, on Catalina Island, was converted in 1978. Across the network the change took nearly sixty years. In a city that switched, the demand for operators fell sharply around the conversion.[1,2]

The people who paid were the operators already in post, particularly older ones. The study found exits from work and moves into lower-paying occupations among affected workers. Younger cohorts entering the labour market largely found other work, with local conditions affecting the adjustment. A rapid change within one city and a gradual change across the network can be true at the same time.[2]

AI reaches some tasks through which newcomers learn a profession: first drafts, document review and basic analysis. If employers take away the tasks without designing alternatives, trainees may lose practice and feedback, even where junior employment holds up. Recent US evidence on early-career employment is suggestive. It cannot by itself establish the size or cause of that training effect. I looked at the mechanism in *Who Will Become the Experts?*, starting with The Disappearing First Rung.[13]

A machine that acts

Earlier systems could already make decisions within defined rules, including telephone switching. A system that uses an AI model to choose among tools can act across a wider, less predictable range. In April 2026 the founder of PocketOS reported that a coding agent had deleted the company's production database and associated backups. The details come from incident accounts and should be treated accordingly. The more common risk may be a plausible action that does the wrong thing and is hard to notice.[3]

A person asked to approve hundreds of routine decisions may become less attentive, particularly if a system is usually correct. Kroese describes asking leading models for numbers from an IMF report and getting plausible figures that did not match it. Human review can work when the reviewer has the

original evidence, time to check it and authority to intervene. Audit, limited permissions and reversible actions can reduce how much rests on the review alone.[4,5]

Geek Box · The developers who felt faster

In early 2025 METR studied 16 experienced open-source developers completing 246 tasks in repositories they knew. It randomly assigned tasks to allow or disallow AI tools. The developers expected a 24% speed-up and afterwards estimated a 20% speed-up. On the study's measure, the tasks for which AI was allowed took 19% longer. The result applies to that sample, those tools and that work. A later METR update explains why its newer estimates should be interpreted cautiously. The IEA's April 2026 report, Table 11.1, cites the original study with the direction reversed. Readers should take the figure from the primary study.[6,14,15]

Working on judgement itself

Computers and previous forms of automation have long reached clerical and analytical work. What may be different here is the breadth of the tasks a general-purpose AI system can attempt, including some used to develop professional judgement. History gives us examples of replacement and adaptation, but no straightforward base rate for whether new routes into expertise will be built soon enough.

There are candidates. Narayanan and Kapoor discuss the work of supervising and controlling AI systems. Perez points to services delivered face to face. Nordhaus's analysis of a possible technological 'singularity', discussed by Kishtainy, also raises a practical point: progress in one part of the economy need not remove bottlenecks in the rest.[7,8,9]

Peña argues that curiosity, critical thinking and self-control become more valuable as AI offers easy answers. Curiosity gives us questions the system has not framed. Critical thinking helps us evaluate the answer and the evidence behind it. Self-control helps us decide when to stop asking and do the work. Whether schools and workplaces give people practice in exercising those skills is a matter of design.[10]

A few firms and their governments

Infrastructure can be concentrated even when the services built on it are widely used. The IEA reports that a small group of cloud and technology companies accounts for a large share of investment in data-centre capacity. The cost of building at the frontier gives governments reasons to pay attention to procurement, access to computing and competition. Nuclear history is useful for thinking about strategic rivalry, though the technologies and the barriers to diffusion differ.[11,12]

The Nuclear Non-Proliferation Treaty entered into force in 1970. Its history illustrates how much cooperation can depend on competing states seeing an interest in a common arrangement. It also warns against confidence about technological leads. RAND recounts General Leslie Groves's 1946 estimate that the Soviet Union was about twenty years from a bomb. It tested one in 1949. For AI, a scientific panel could help establish evidence without having the authority to enforce controls.[12]

History cannot tell us how capable these systems will become. It can tell us who tends to gain from a new technology, and what decides it. That is Part 5.

For a related account of why a fluent answer can be difficult to check, see Why AI Makes Things Up.

Continue the series

Previous · Part 3 · How institutions respond

Next · Part 5 · Who gains, and what to expect

Sources and notes

1. Feigenbaum, J. J. and Gross, D. P., 'Organizational and Economic Obstacles to Automation: A Cautionary Tale from AT&T in the Twentieth Century', NBER Working Paper 29580, as revised. Source
2. Feigenbaum, J. and Gross, D. (2024), 'Answering the Call of Automation: How the Labor Market Adjusted to Mechanizing Telephone Operation', Quarterly Journal of Economics 139(3), pp. 1879–1939. Source
3. Euronews (28 April 2026), reporting PocketOS founder Jer Crane's account of the April 2026 database incident. The factual detail is attributed to the company, and independent verification of its full account is unavailable here. Source
4. Narayanan, A. and Kapoor, S. (2025), AI as Normal Technology, Knight First Amendment Institute, pp. 16–17, citing Parasuraman, R. and Manzey, D. H. (2010), 'Complacency and Bias in Human Use of Automation', Human Factors 52(3), pp. 381–410. Source
5. Kroese, B. (2025), 'Instilling Trust through Statistics', Finance & Development, December 2025, p. 10. Source
6. Becker, J., Rush, N., Barnes, B. and Rein, D. (2025), Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity, METR, July 2025. Source
7. Narayanan and Kapoor (2025), pp. 17–18. Source
8. Perez, C. (2010), The Financial Crisis and the Future of Innovation, Technology Governance Working Paper 28, pp. 19–20. Source
9. Kishtainy, N. (2025), 'A New Industrial Revolution?', Finance & Development, December 2025, pp. 46–49. Source
10. Peña, P. A. (2025), 'The Human Edge', Finance & Development, December 2025, pp. 50–53. Source
11. International Energy Agency (2026), Key Questions on Energy and AI, sections on technology company investment and data-centre capacity. Source
12. Boudreaux, B., Smith, G., Geist, E. and Dion, L. (2025), Insights from Nuclear History for AI Governance, RAND, PE-A3652-1. Source
13. Brynjolfsson, E., Chandar, B. and Chen, R. (2025), Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, working paper, November 2025 version. Source
14. METR (24 February 2026), We are Changing our Developer Productivity Experiment Design. Source
15. International Energy Agency (2026), Key Questions on Energy and AI, Table 11.1. The supplied April 2026 edition reverses the direction of the 2025 METR finding. Source