WHAT CHANGED SINCE SEPTEMBER?

AI is making discoveries. We still have to steer it.

Since September, scientific discovery looks stronger and security failures look more widespread. Reliable replacement of whole research teams remains unproven.

AI 2027 describes a possible future in which AI rapidly improves itself and becomes far more capable than people. This review checks that story against the evidence.

FIVE AREAS OF PROGRESS	LAST UPDATE	LATEST EVIDENCE	WHAT IT MEANS
01 **What AI can do**	Strong results in coding, reasoning and selected professional tasks.	OpenAI and Anthropic report mathematical and biological discoveries by large groups of AI agents, with human guidance. [1,2]	AI can contribute more to discovery. Success on one problem does not establish expertise everywhere.
02 **Working independently**	Defined tasks were easier to delegate than long, uncertain projects.	Ongoing AI workers are becoming available. Large groups tackle difficult projects, with people still directing and checking them. [1,2,3,4]	More work can be delegated. Dependable everyday performance still needs stronger evidence.
03 **Improving AI itself**	AI helped researchers write code and run experiments.	More research tasks are automated. Evidence that AI is dramatically accelerating its own development remains limited. [3]	New scientific discoveries and faster AI development are different achievements. A runaway self-improvement cycle remains unproven.
04 **Jobs and the economy**	Investment and adoption were ahead of clear benefits across the economy.	Capable AI is cheaper. A Texas study strengthens evidence of employment pressure on some graduates. [5,6]	Disruption can affect particular workers before it becomes obvious in national figures.
05 **Safety and control**	Serious incidents showed that powerful AI could cross security boundaries.	Australia confirmed an AI intrusion into a Medicare statistics website during routine research. Some advanced AI research is paused. [7,8,9]	An ordinary assignment can lead beyond permission given. Governments are responding; reliable limits and shutdowns matter.

An AI agent can use tools and take actions, rather than only answer questions.

THE IMPORTANT CONNECTION

An AI system can get better at producing a result without becoming better at following the rules.

Where might this lead?

ESTIMATES UNCHANGED THIS MONTH

These dates require dependable performance across almost all the relevant work. A striking result on one problem is not enough to move them.

AROUND

2029

AI takes over coding

A leading AI lab could rely on AI for its full range of software-engineering work.

AROUND

2033

AI matches top experts

Across almost all work that requires thinking, learning and judgement.

AROUND

2037

AI goes far beyond people

Substantially surpasses the best humans across almost all of that intellectual work.

This review's uncertain estimates of what will happen in the world. Unchanged this month; separate from the original AI 2027 story and its creators' forecasts. Earlier or later arrival remains possible.

Evidence limits. Scientific results are lab-reported. Research-acceleration estimates are preliminary. Labour studies cannot establish every cause of weaker hiring. The Australian incident occurred in June and was disclosed in September; patient-record access was not established.

Evidence available at 1 October 2026, 09:27:32 AEST. The table distinguishes evidence and interpretation; the dates are forecasts.

1 OpenAI discovery | 2 Anthropic discovery | 3 METR research assessment | 4 OpenAI ongoing agents | 5 GPT-6.1 Sol | 6 Dallas Fed labour study | 7 Australian Government | 8 OpenAI containment failure | 9 OpenAI safety proposals