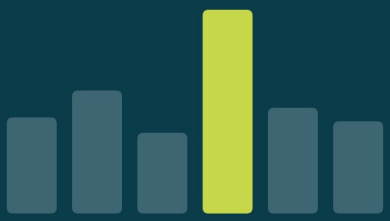
COMPANION READ

The gist of AI

What it is, where it came from, and how to think about it

1955: one summer. 2026: the redesign.



Jaime Díaz-Beltrán

Data science and AI leader · October 2026

Facts are current as of 29 October 2026. AI moves fast: check the date on a figure before you quote it.

BEFORE YOU START

How to read this document

This document follows the talk act by act, with the detail, evidence and caveats a slide cannot hold. Pick the depth you have time for.

5 min	20 min	45 min
Read Chapter 1: the whole talk on one page.	Read the “In short” box and the figures in each chapter.	Read it all, with the evidence and the caveats.

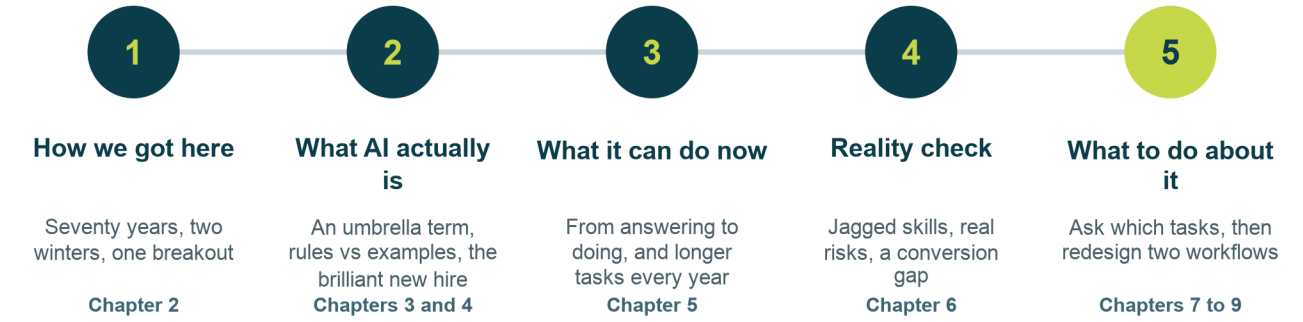
Every number has a source in brackets, numbered as in the list at the end. Where a figure comes from a company that sells AI, I say so. Where the evidence is thin or disputed, I say that too.

What’s inside

1	**The talk on one page**	The whole argument in five acts
2	**How we got here**	Seventy years, two winters, one breakout
3	**What AI actually is**	An umbrella term, rules vs examples, the brilliant new hire
4	**Key terms**	Fifteen terms in plain language, each with a picture
5	**What AI can do now**	Four ways to work with it, agents, and longer tasks every year
6	**Reality check**	The jagged frontier, four risks, the conversion gap, the electricity lesson
7	**What it means for professional services**	Tasks not jobs, the task map, the pyramid
8	**Mental models to keep**	Eleven pictures that hold up under pressure
9	**A 90-day plan**	Two redesigns beat twenty pilots
10	**Where to go deeper**	Books, courses and videos, with Spanish editions
11	**Sources**	Every reference, numbered as cited, plus image credits

CHAPTER 1

The talk on one page



1 HOW WE GOT HERE

- **AI is old.** The question dates to 1950 and the name to 1955. What changed in 2022 was access: anyone could type to it.

2 WHAT AI ACTUALLY IS

- **"AI" is an umbrella.** Your pricing, fraud and churn models are AI too. The new part is generative AI, and large language models (LLMs) in particular.
- **Regular software follows rules. AI learns from examples.** That is why it handles messy inputs, and why it can be fluent, confident and wrong.
- **Think of an LLM as a brilliant new hire:** it has read the whole library, has never met your client, and rarely says "I don't know."

3 WHAT IT CAN DO NOW

- **AI has moved from answering to doing.** You can ask it, work alongside it, delegate to an agent, or run agents that keep working after you leave. You choose how far.
- **The tasks it can finish are getting longer, fast.** On software tasks, from about 4 minutes (2023) to about 12 hours (early 2026) at 50% success. At 80% success, about an hour.

4 REALITY CHECK

- **Its skills are jagged.** It can score at gold-medal level at the Math Olympiad and misread a clock half the time. Test it on your own work to find the edge.
- **Almost every firm uses AI. Few see it in the numbers yet.** Call it a conversion gap: time saved by people is not yet becoming firm results. Checking the work is the new bottleneck.
- **We have seen this before.** Electric motors paid off only when factories redesigned the floor. Most AI today is a new motor on an old floor.

5 WHAT TO DO ABOUT IT

- **Ask which tasks, not which jobs.** But be honest about the second step: when enough tasks move, jobs change, and the first place it shows is junior hiring.
- **Two redesigns beat twenty pilots.** Use it yourself every week, redesign two or three workflows in 90 days, and agree on the outcome before anyone builds.

CHAPTER 2

How we got here

Seventy years, two winters, one breakout.

IN SHORT

- The question dates to 1950 and the name to 1955. This is not a new field.
- Twice, “AI winters” followed promises that ran far ahead of results.
- What changed in 2022 was access: for the first time, anyone could type to it.

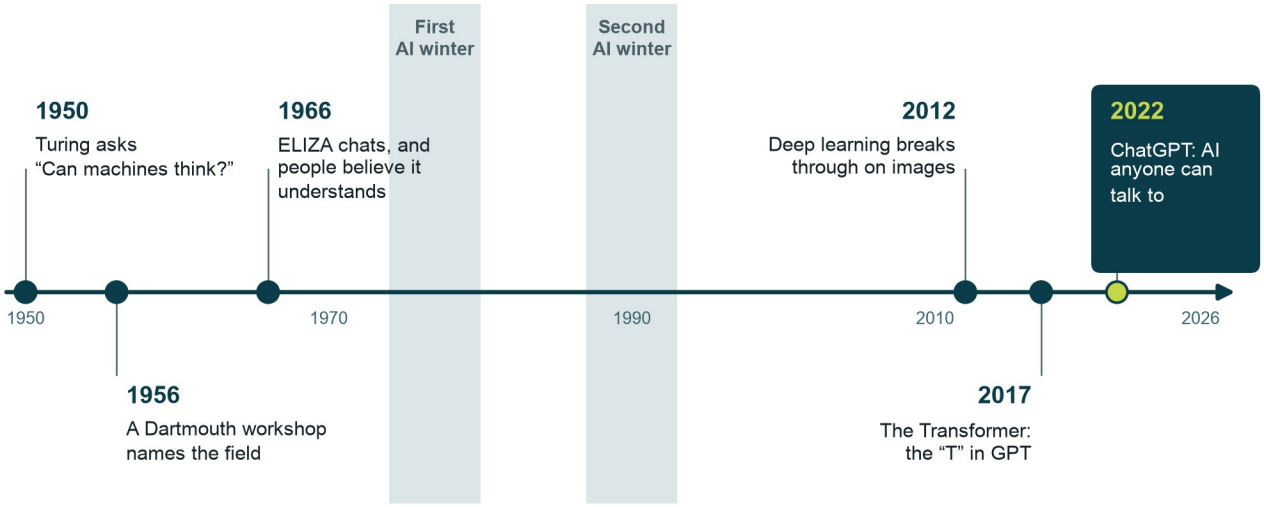


Figure 1 Seventy years of AI in six moments. Shaded: the two AI winters, when funding dried up. [1, 2, 3, 4, 5]

A seventy-year-old idea

In October 1950 Alan Turing opened a paper with the question “Can machines think?” and proposed replacing it with a test he called the imitation game. [6]

Five years later, on 31 August 1955, John McCarthy, Marvin Minsky, Nathaniel Rochester and Claude Shannon asked the Rockefeller Foundation to fund a summer workshop at Dartmouth College. Their proposal is where the term “artificial intelligence” comes from. [7]

> *“We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire.”*
>
> McCarthy, Minsky, Rochester and Shannon (1955)

They thought a summer would get them a long way. It took seventy years.

The 1960s brought the first boom. In 1966 Joseph Weizenbaum’s ELIZA, a simple pattern-matching chatbot, convinced some users it understood them. His warning still applies to today’s chatbots:

> *“ELIZA shows, if nothing else, how easy it is to create and maintain the illusion of understanding, hence perhaps of judgment deserving of credibility. A certain danger lurks there.”*
>
> Joseph Weizenbaum, Communications of the ACM (1966)

Then came two “AI winters”. The first followed the UK’s Lighthill report in 1973 (“In no part of the field have the discoveries made so far produced the major impact that was then promised”). The second came in the late 1980s, when the market for expert systems and specialised AI hardware collapsed. [8]

The current wave started in 2012, when a deep neural network called AlexNet won the ImageNet image-recognition contest by a wide margin (15.3% error against 26.2% for the runner-up). In 2017 Google researchers published the Transformer, the architecture behind almost every model since: it is the “T” in GPT.

Milestones

Each of these felt like the moment machines became intelligent. Each one, soon after, felt like ordinary software.

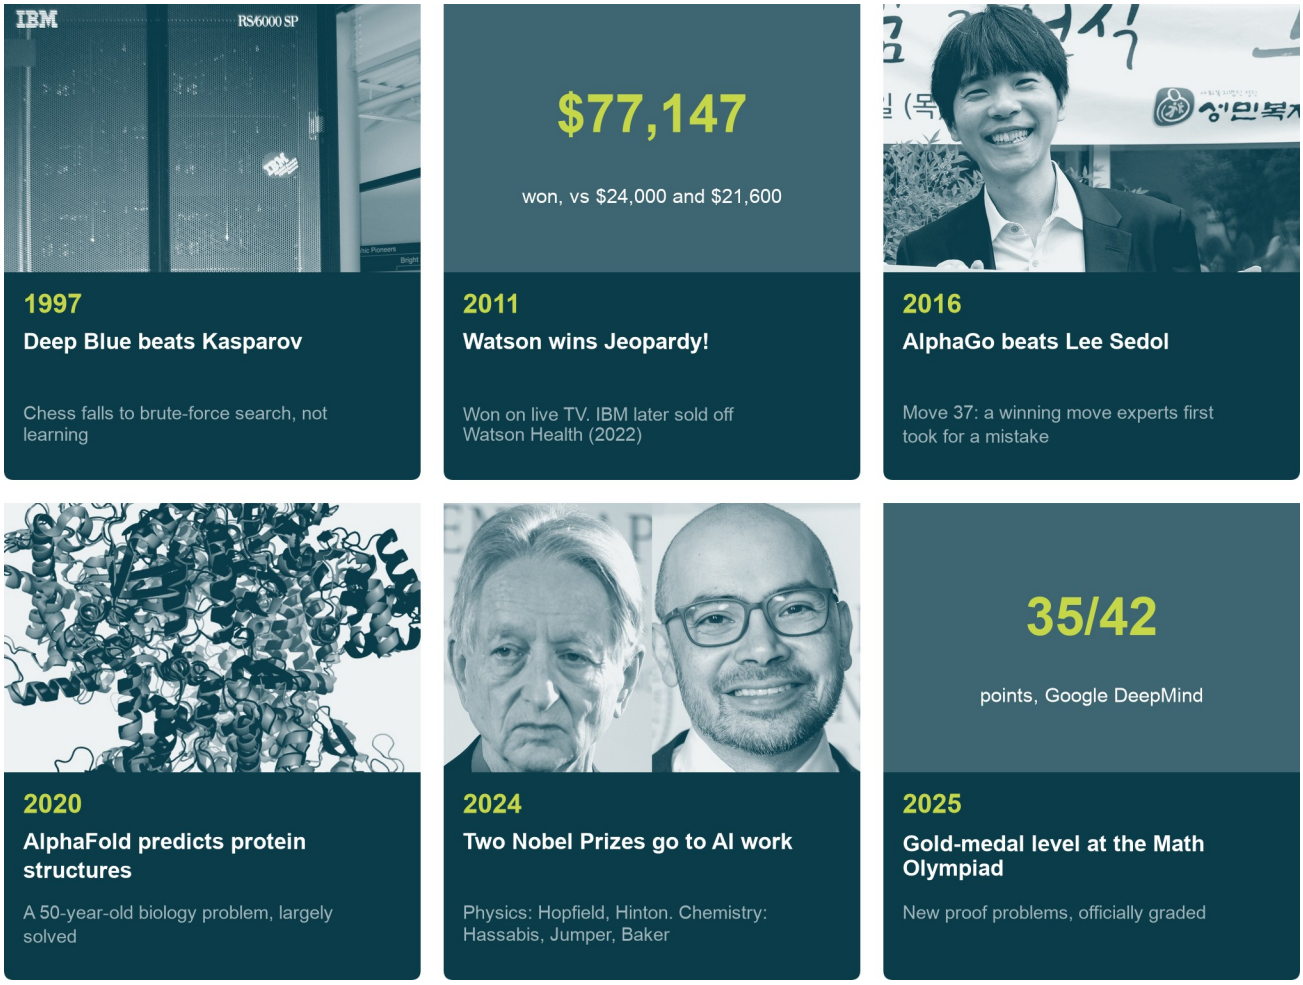


Figure 2 Six landmarks, from winning our games to advancing science. Photos: Erik Pitti, CC BY 2.0; LG Electronics, CC BY 2.0 (cropped); Jumper et al., Nature 2021, CC BY 4.0 (cropped); John Sears, CC BY-SA 4.0 (cropped). All via Wikimedia Commons, toned. [9, 10, 11, 12, 13, 14]

THE DETAILS

- Deep Blue won 3.5 to 2.5 in May 1997, the first match win over a reigning world champion. It won by brute-force search, not learning.
- Watson won $77,147 on Jeopardy! against $24,000 and $21,600. IBM later sold off Watson Health: a win on stage is not value in a business.
- AlphaGo's move 37 in game 2, which experts first took for a mistake, had about a 1 in 10,000 chance of being played by a human.
- AlphaFold 2 reached near-experimental accuracy at CASP14 in 2020. By 2022 its database covered over 200 million protein structures.
- The 2024 Nobel Prizes: Physics to Hopfield and Hinton; Chemistry to Hassabis and Jumper, shared with Baker.
- At the 2025 Math Olympiad, Gemini Deep Think scored 35 of 42, officially graded. No AI got a medal: medals are for students.

Why 2022 felt like the beginning

ChatGPT was not dramatically more powerful than what the labs already had. What changed was access. Bharat Anand and Andy Wu of Harvard Business School compare it to the move from MS-DOS to Windows, which transformed computing "not by making computers significantly more powerful but by allowing people to access that power without knowing MS-DOS commands." [15]

~100M

monthly ChatGPT users two months after launch. A UBS estimate, not an OpenAI figure [16]

1.2B

weekly users by September 2026, as self-reported by OpenAI [17]

2 in 3

S&P 500 earnings calls mentioned "AI" in 2026 [18]

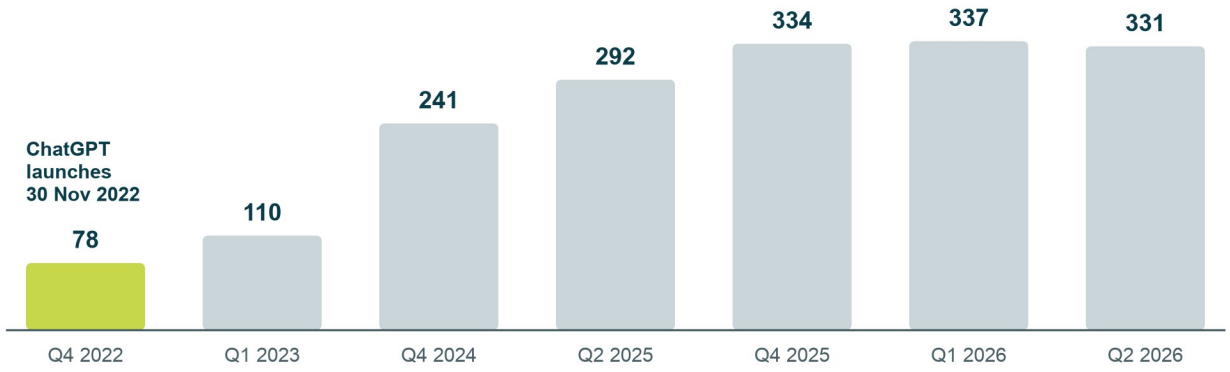


Figure 3 S&P 500 companies citing "AI" on earnings calls. FactSet counted companies until 2024 and calls from 2025. [18]

Collins Dictionary named "AI" its Word of the Year for 2023. Cambridge chose "hallucinate". [19]

The AI effect

> *"Intelligence is whatever machines haven't done yet."*
>
> Larry Tesler, computer scientist, in his own wording of what became "Tesler's theorem" [20]

Spam filters, credit scoring, route optimisation and speech recognition were all "AI" once. Today they are just software. Expect the same to happen to chatbots.

CHAPTER 3

What AI actually is

An umbrella term, a new kind of software, and one picture to hold it all.

IN SHORT

- "AI" covers very different things. When someone says AI, ask which kind.
- Regular software follows rules people write. AI learns patterns from examples, so it handles messy inputs, and it can be fluent, confident and wrong.
- A large language model is the most well-read autocomplete ever built. Think of it as a brilliant new hire.

An umbrella, not one thing

Arvind Narayanan and Sayash Kapoor of Princeton put it best: "AI" is an umbrella term for a set of loosely related technologies. Imagine having only the word "vehicle" for bikes, cars, buses and spacecraft, and then arguing about whether vehicles are good for the environment. [21]

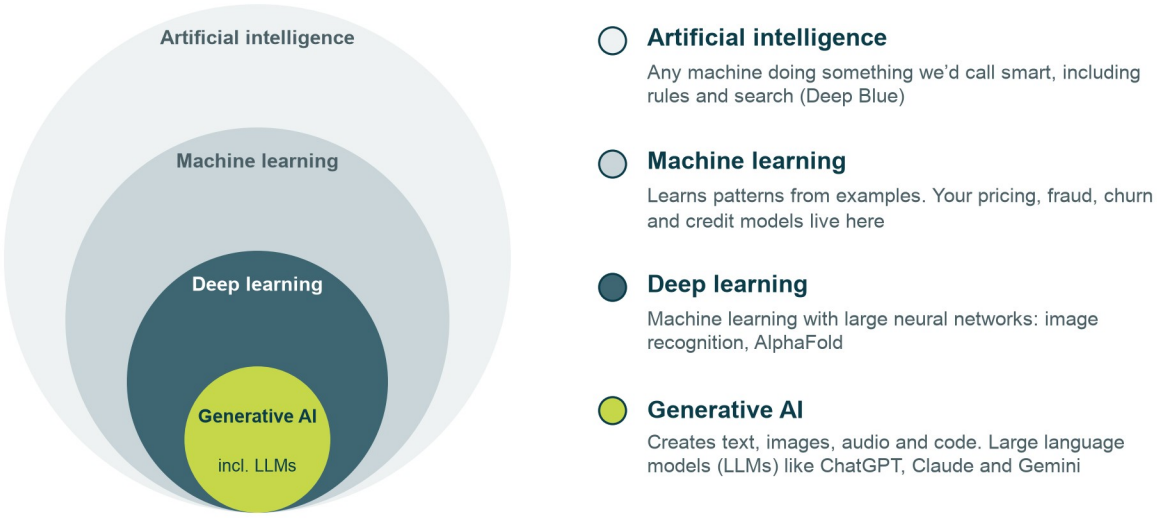


Figure 4 How the terms nest. You have had AI in your pricing and fraud models for years; the new part is generative AI.

The nesting is a simplification. Generative AI has no formal definition and sits mostly, not entirely, inside deep learning. Agents, retrieval and "reasoning" are not smaller circles but systems built around a language model.

> **Think of it as** the word “vehicle”. When someone says AI, ask which kind.

Regular software follows rules. AI learns from examples.

> *“Explain with examples, not instructions.”*
>
> Cassie Kozyrkov, former Chief Decision Scientist, Google (2018)

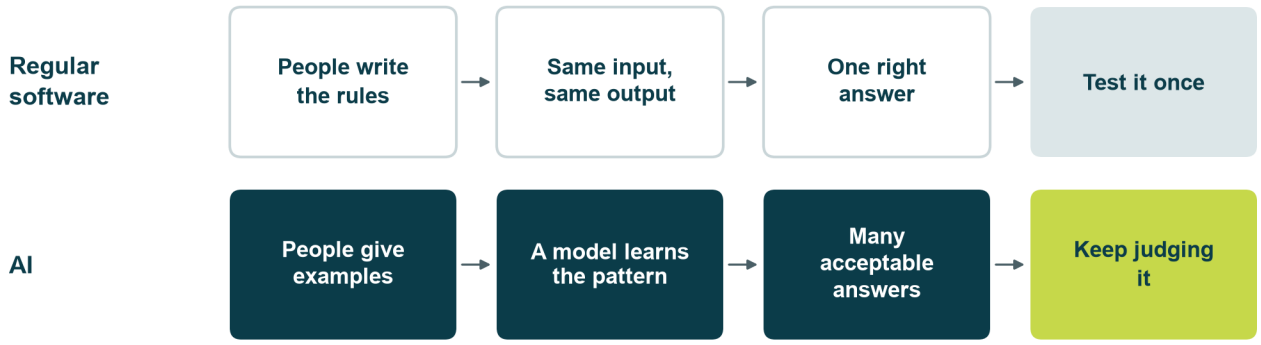


Figure 5 Two ways to build software. [22, 23]

Regular software is a set of rules people write: the same input gives the same output, there is one right answer, and you can test it once. Machine learning flips this. You show the system thousands of examples and it learns the pattern, which is why it can handle inputs no one could write rules for: an email, a contract, a meeting transcript.

The price is that the output is probabilistic. Ask twice and you may get two different answers, both acceptable. With many acceptable answers, you cannot test the system once and forget it. You have to keep judging it, which is why measuring value with AI looks different from measuring it with regular software.

> **Think of it as** a rulebook you write, versus a recipe the computer writes from thousands of your examples.

How a large language model works

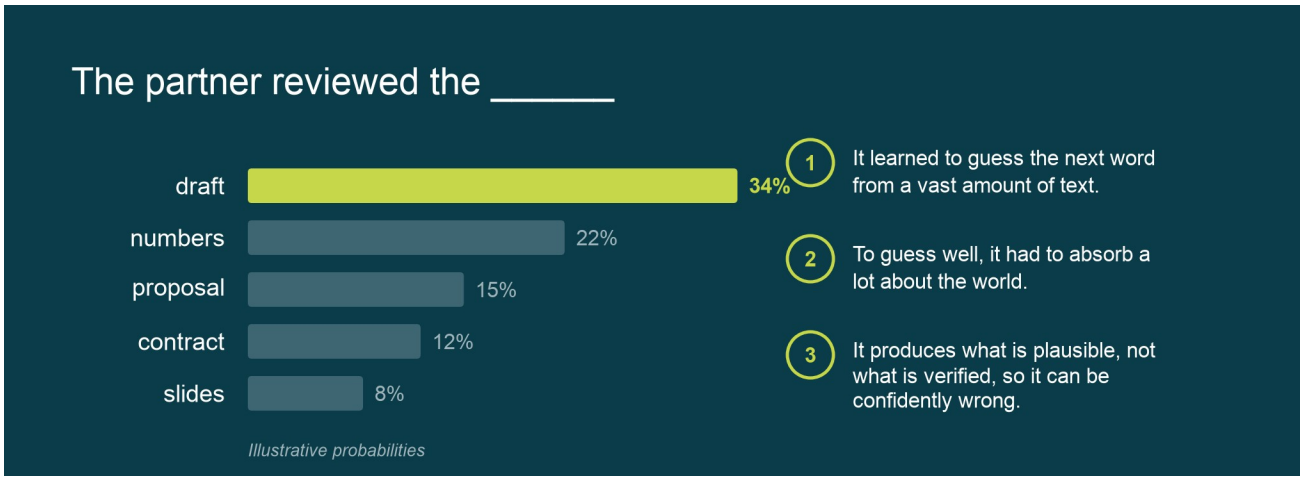


Figure 6 A language model is trained to guess the next word. The probabilities shown are illustrative. [24, 25]

A large language model is trained as autocomplete. It reads a vast slice of the internet and learns to predict the next word, billions of times over. To predict well, it has to absorb a great deal about how the world works. As 3Blue1Brown puts it, it is “a sophisticated mathematical function that predicts what word comes next”.

Two consequences follow:

- **It produces what is plausible, not what is verified.** That is why it can invent a source with total confidence.
- **Next-word prediction is how it is trained, not the ceiling of what it can do.** Newer “reasoning” models work through intermediate steps before answering, and research shows models plan several words ahead.

The brilliant new hire

The picture used through the talk: an LLM is like a brilliant new hire. It has read the whole library, it is fast and tireless, it has never met your client, and it rarely says “I don’t know.” Most AI jargon maps onto that picture, as the next chapter shows.

Where the picture breaks, and it does:

- A person learns from your red pen. The model does not, unless someone builds that feedback in.
- A person forgets slowly. The model forgets the conversation when it ends, unless it is given memory.
- A person is accountable and covered by your professional indemnity. The model is neither.

CHAPTER 4

Key terms

Fifteen terms that cover most AI conversations, each with a plain definition and a picture to remember it by.

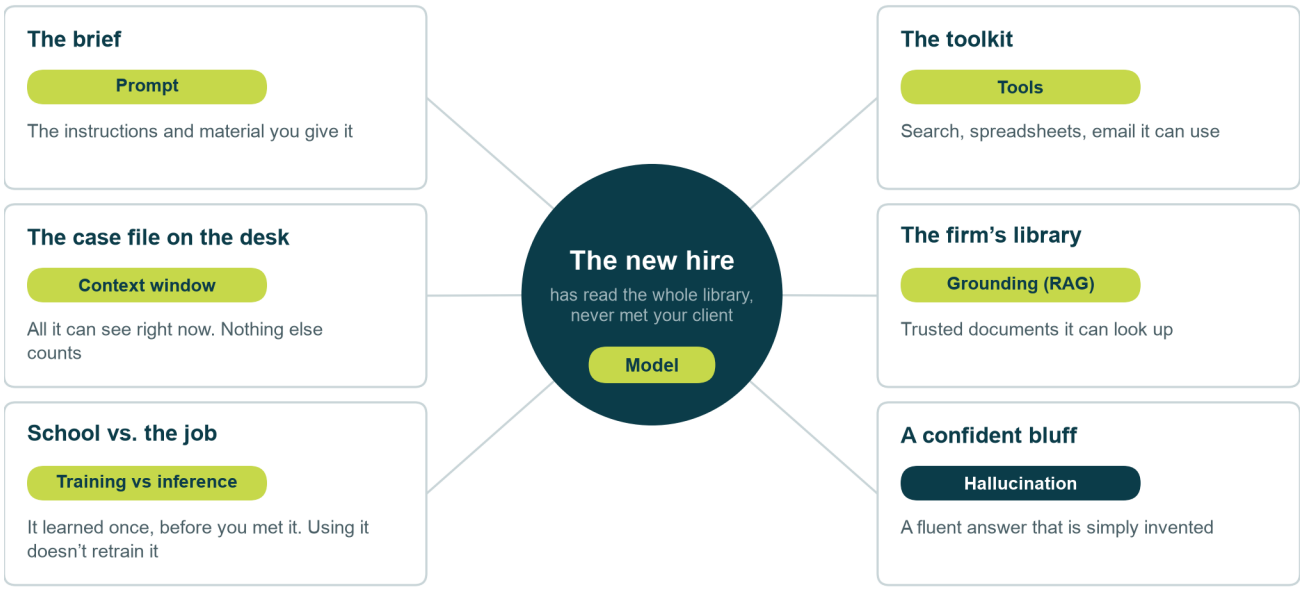


Figure 7 Most AI jargon maps onto something you already manage: a brilliant new hire.

Definitions draw on the Google Machine Learning Glossary, Anthropic's glossary, NIST and the OECD. [26, 27, 28, 29]

AI system

A machine-based system that infers from its inputs how to produce predictions, content, recommendations or decisions (OECD, 2023).

Think of it as *the word "vehicle".*

Machine learning

Software that learns patterns from examples instead of following hand-written rules.

Think of it as *training a new analyst on 10,000 past cases instead of a rulebook.*

Model

The trained system itself: the patterns learned from data.

Think of it as *the brilliant new hire.*

Neural network and deep learning

A model made of many layers of numeric "dials" tuned on data. Deep means many layers.

Think of it as *a mixing desk with billions of knobs.*

Generative AI

AI that creates new text, images, audio or code.

Think of it as *a chef who cooks a new dish, not a critic who grades it.*

Large language model (LLM)

A generative model trained on vast amounts of text to predict the next word.

Think of it as *the most well-read autocomplete ever built.*

Training vs inference

Training builds the model once, at great cost. Inference is every time you use it. Using it does not retrain it.

Think of it as *medical school vs. a doctor's appointment.*

Parameters

The learned numbers inside a model. Their count signals size, not quality.

Think of it as *the knobs on the mixing desk.*

Token

A chunk of text, often part of a word. The unit AI is priced and limited in.

Think of it as *the meter on a taxi.*

Prompt

The instructions and material you give the model.

Think of it as *the brief to a new associate.*

Context window

Everything the model can see at once: prompt, files and the conversation so far.

Think of it as *the case file on the desk.*

Hallucination

Fluent, confident output that is false. NIST calls it confabulation.

Think of it as *a confident bluff.*

Grounding (RAG)

Giving the model trusted documents at question time so answers rest on sources. It reduces made-up answers; it does not eliminate them.

Think of it as *an open-book exam.*

Reasoning model

A model that works through steps before answering: better on hard problems, slower and costlier.

Think of it as *showing your work on scratch paper.*

Agent

A model that uses tools in a loop to reach a goal, with less supervision.

Think of it as *the new hire who takes a workstream and runs with it.*

THE ONE TO REMEMBER

Hallucination is a feature of how it works

It predicts what is plausible. Grounding and review reduce invented content; nothing removes it entirely.

CHAPTER 5

What AI can do now

From answering questions to doing the work, and for longer stretches every year.

IN SHORT

- There are four ways to work with AI: ask, assist, delegate and always-on. You choose how far, task by task.
- An agent is a model that uses tools in a loop until a goal is met. Orchestration manages many of them.
- The length of task AI can finish has doubled about every four months since 2023. Reach is growing faster than reliability.

Four ways to work with AI

All four exist today. Higher is not always better; it depends on the task and the cost of a mistake. Andrej Karpathy calls this the autonomy slider: you decide, task by task, how much the AI does alone. [30]

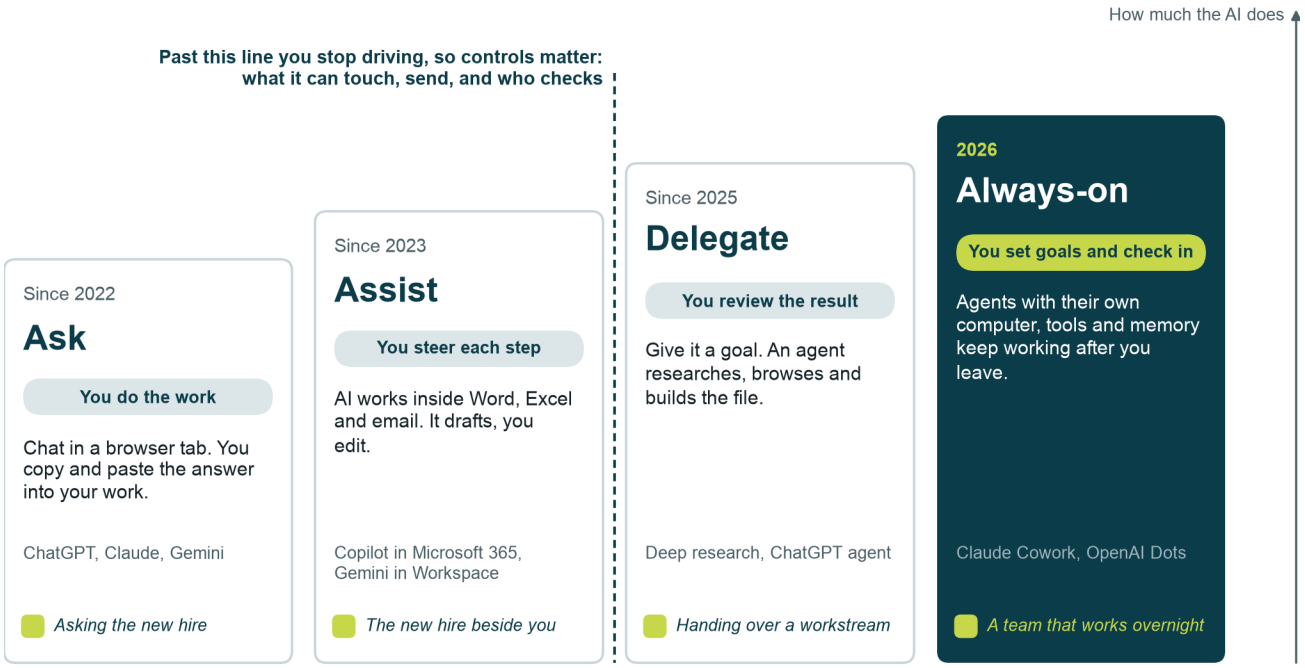


Figure 8 Four ways to work with AI, all available today. [5, 31]

The important line runs between Assist and Delegate. To the left, you are still driving. To the right, you have stopped driving, and you need controls: what the agent may touch, what it may send, and who checks before anything reaches a client.

WORTH KNOWING

- Always-on agents are new and moving fast. Claude Cowork and OpenAI Dots launched in 2026. Check what is available in Colombia, and under your firm's data rules, before relying on them.
- Two others get attention: OpenClaw (open source, with documented security incidents) and Meta's Muse (US only).

Agents and orchestration

The classic textbook definition, from Russell and Norvig, is broad: "An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators." Today the word usually means something narrower: a language model that uses tools in a loop to reach a goal. It plans, uses a tool, checks the result and goes again until it is done. [32]

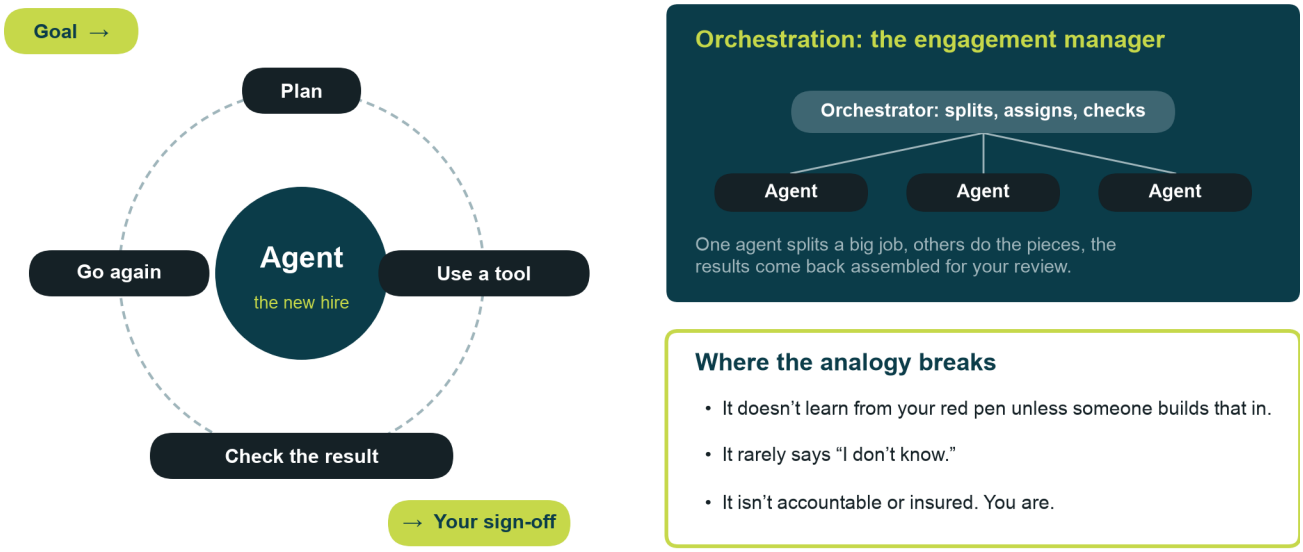


Figure 9 An agent works in a loop until the goal is met and you sign off. Orchestration splits a big job across several agents. [33]

Anthropic draws a useful line between workflows, where the steps are fixed in advance, and agents, which decide their own steps. Its advice is to use the simplest option that works. Orchestration is the next layer: one agent splits a large job, hands the pieces to other agents, checks what comes back and assembles it, much like an engagement manager.

How long a task AI can handle

METR, an independent research group, measures how long a task AI agents can complete, in the time it takes a skilled professional. A model's "time horizon" is the length at which it succeeds half the time. [34, 35]

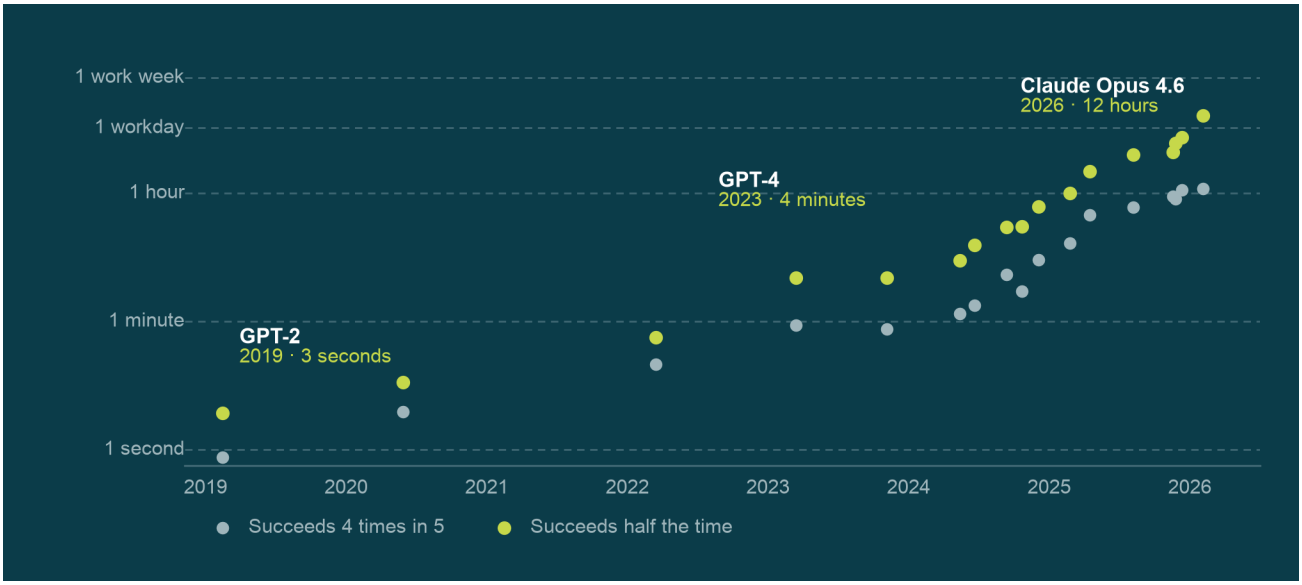


Figure 10 Task length that public frontier models complete, measured by how long the task takes a skilled professional. Log scale: each ruler is a big step up.

~12 h
of expert work the best public model finishes half the time (Claude Opus 4.6, February 2026)

~70 min
for the same model at 80% success. Reliability lags reach

~4 mo
for that task length to double since 2023 (about 7 months over 2019 to 2025)

> **Think of it as** a new hire who can attempt a full day's job, but whom you would leave alone for about an hour.

WORTH KNOWING

- The tasks are software, machine-learning and cybersecurity work with clear success criteria. Other kinds of work, such as operating a computer screen, show much shorter horizons.
- A 50% horizon is a coin flip. The 80% horizon is five to ten times shorter.
- METR stopped updating its page in September 2026, and says measurements above 16 hours are unreliable with its current tasks. Exponential trends eventually bend; nobody knows when.

CHAPTER 6

Reality check

Jagged skills, real risks, and a gap between using AI and getting value from it.

IN SHORT

- AI's skills are jagged: brilliant at some hard tasks, unreliable at some easy ones. You find the edge only by testing it on your own work.
- Four risks stay with you: client data, invented content, regulation and over-trust.
- Almost every firm uses AI, yet few see it in the numbers. The old process, not the technology, is the constraint.

The jagged frontier

The best evidence for consultants comes from consultants. In a randomized trial with 758 BCG consultants, now peer-reviewed in Organization Science (2026), researchers from Harvard, Wharton and BCG gave some participants GPT-4. [36]

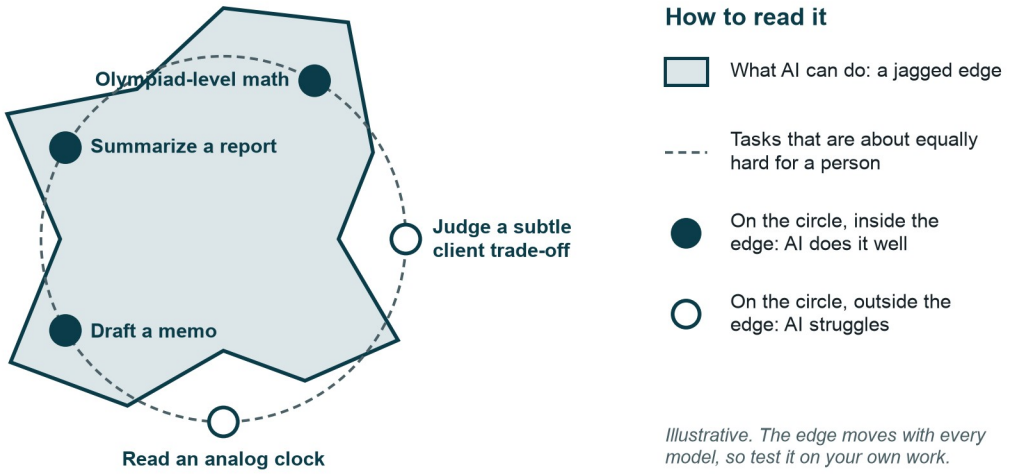


Figure 11 AI's strengths don't follow human difficulty. Models that score at gold-medal level at the Math Olympiad read analog clocks correctly only about half the time. [37]

Inside the frontier, consultants with AI did better on every measure:

+12%	25%	~⅓
more tasks finished	faster	higher quality (about 32% in the published paper; the 2023 working paper said 40%)

Just outside it, consultants with AI were right 19 percentage points less often: about 85% correct without AI, 60% to 71% with it. Their wrong answers were rated more polished and persuasive.

A follow-up study found three working styles: “centaurs”, who split tasks between themselves and the AI and were the most accurate; “cyborgs”, who blend the two; and “self-automators” (27%), whose work was fast and polished but lacked depth. [38]

Fair objection: that trial used GPT-4 in 2023. The frontier has moved. It is still jagged, and you still cannot see the edge until you test it on your own work.

Four risks you can’t delegate

The new hire works fast. The accountability stays with you.

Client data

The paid business plans of ChatGPT, Claude, Microsoft Copilot and Gemini all state in writing that they don’t train on customer data by default. Personal accounts are different: ChatGPT’s free and personal plans share data for training unless the user opts out. Shadow use is common: 78% of people who use AI at work bring their own tools, and 44% of employees in a 47-country study admit using AI in ways that break policy. [39]

Invented content

Deloitte refunded A$97,587 of an A$440,000 report for Australia’s employment department after it was found to cite works that don’t exist and a fabricated court quote (2025). In New York, a judge fined the lawyers suing Avianca $5,000 for citing cases ChatGPT made up (2023). In Colombia, the Constitutional Court reviewed a judge’s use of ChatGPT in 2024 and held that AI cannot replace the judge. [40, 41]

Regulation

The EU AI Act lists AI used in recruitment and managing workers, and in “risk assessment and pricing in relation to natural persons in the case of life and health insurance”, as high-risk. After the 2026 Digital Omnibus, those duties apply from 2 December 2027. Colombia has no AI law yet, but Ley 1581 and the SIC’s Circular Externa 002 de 2024 already govern personal data in AI systems, and CONPES 4144 (2025) sets the national AI policy. [42, 43]

Over-trust

In the BCG study, wrong answers produced with AI were rated more polished and persuasive. Someone accountable must read and sign off every client deliverable. [36]

The conversion gap

The best independent data on firm-level results is a 2026 survey of about 9,000 executives in seven countries, run by economists with several central banks. [44]

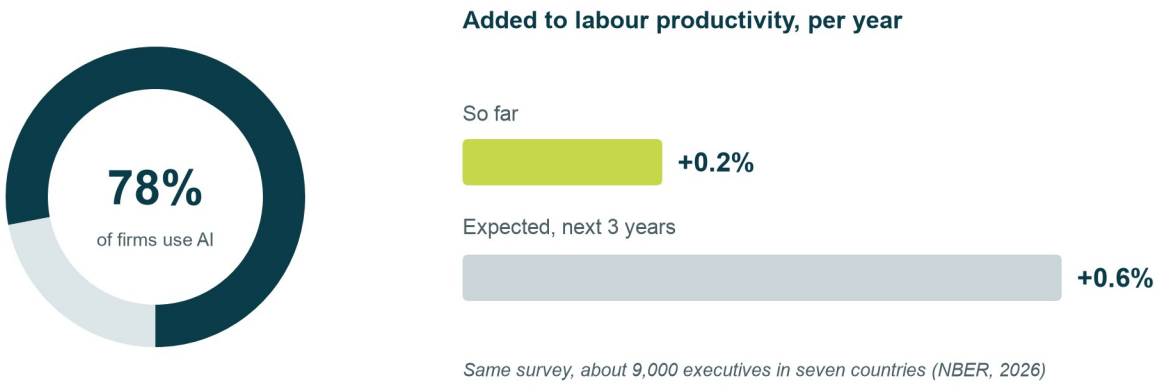


Figure 12 High use, small gains so far, from the same survey.

Other studies point the same way, with caveats worth stating:

- BCG finds the share of companies getting value from AI at scale rose from 5% in 2025 to 7.5% in 2026. BCG sells AI services; read accordingly. [45]
- In a Wharton survey, about three in four large US firms said their AI projects already show positive returns. This is self-reported, and senior leaders were more optimistic than middle managers. [46]
- In Denmark, administrative data on 25,000 workers in exposed occupations found precise zero effects on earnings and hours two years after ChatGPT. [47]

All of these can be true at once. Projects pay off while firm-level numbers barely move.

> **The conversion gap:** people save time, and the firm does not yet turn that time into results. Time saved is not value captured until someone decides what the time is for.

The review tax

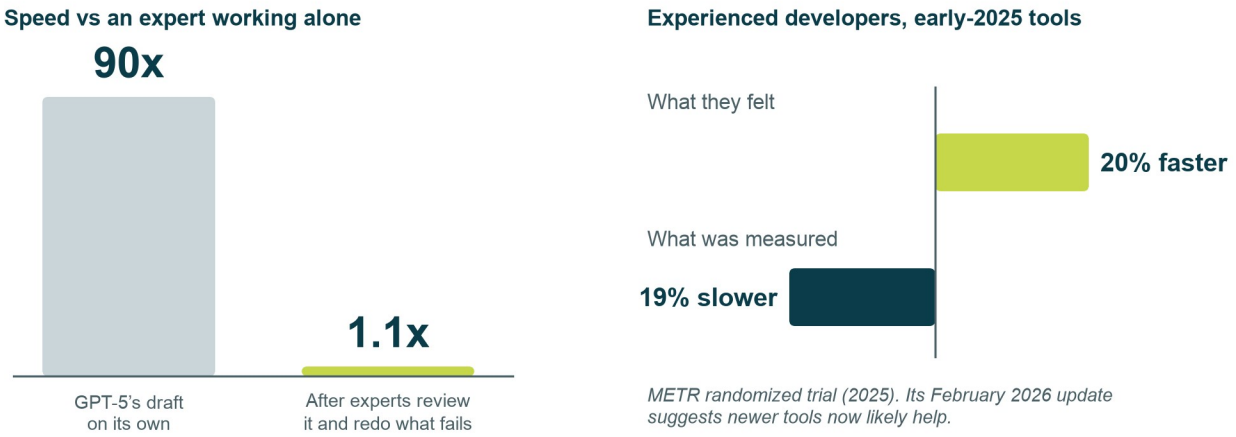


Figure 13 Drafting got cheap. Checking did not, and we are poor judges of our own speed. [48, 49]

In OpenAI's GDPval study, GPT-5 produced expert-level work about 90 times faster than experts working alone. Once an expert reviewed the output and redid from scratch the tasks it failed, the gain was about 1.1 times. That is OpenAI grading its own model, but the peer-reviewed BCG study makes the same point.

In METR's randomized trial with early-2025 tools, experienced developers believed AI made them 20% faster and were measured 19% slower. METR's February 2026 update suggests newer tools now likely speed developers up. The lesson stands: measure, don't trust the feeling.

The electricity lesson

Economists have seen this puzzle before. When electric motors arrived in factories, owners pulled out the steam engine and put a big motor in its place, keeping the same shafts, belts and layout. Productivity barely moved. The payoff came decades later, in the 1920s, when factories put a small motor in every machine and redesigned the floor around the flow of work. Paul David documented this in 1990. [50]

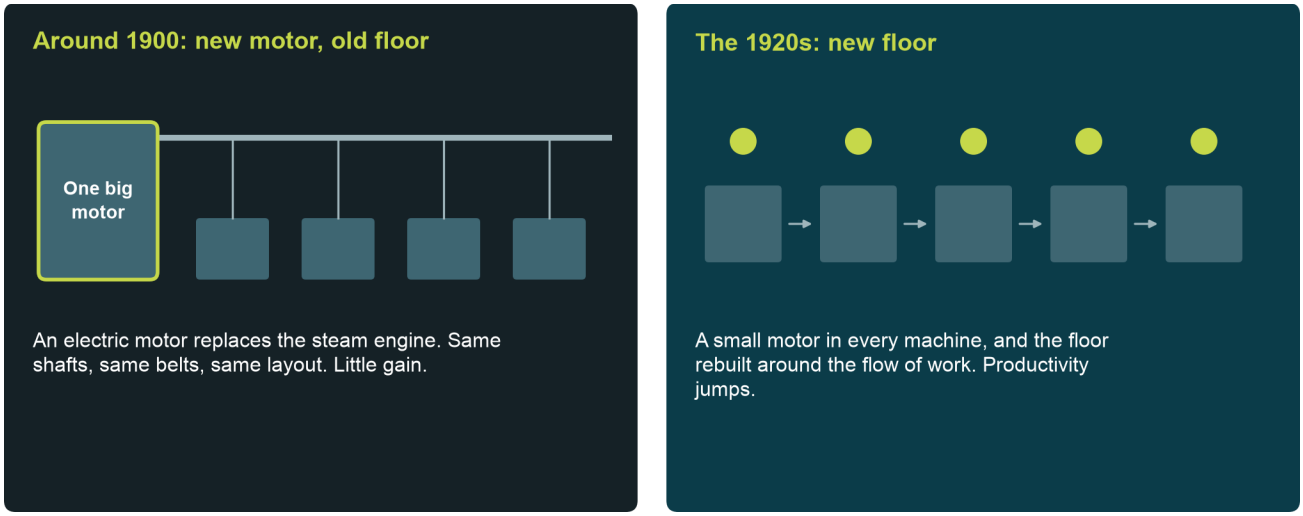


Figure 14 The payoff came from redesigning the floor, not from swapping the motor.

> *"It is time to stop paving the cow paths. Instead of embedding outdated processes in silicon and software, we should obliterate them and start over."*
>
> Michael Hammer, Harvard Business Review, July–August 1990 [51]

Hammer's best-known case was Ford's accounts payable: a redesign achieved a 75% reduction in that department's headcount, against the 20% a conventional improvement program would have delivered. Agrawal, Gans and Goldfarb make the same point for AI: the big gains come from system solutions, not point solutions. [52]

Will it take forty years again? The lag then was the time it took to rebuild factories. Today the rebuild is workflows, roles and pricing, which can move in months. The firms that move first will set the terms for everyone else.

> **Think of it as** most AI today: a new motor on an old factory floor.

CHAPTER 7

What it means for professional services

Ask which tasks, not which jobs, and be honest about what happens when enough tasks move.

IN SHORT

- Jobs are bundles of tasks. AI takes tasks first, and jobs change when enough tasks move.
- A simple map sets AI's role for each task: how costly is a mistake, and does the task run on explicit data or tacit judgment?
- For firms built on leverage, three questions follow: who learns, who pays, and who competes.

Tasks, not jobs, until it takes enough tasks

> *"Everybody wants to talk about job loss, but really, what you want to look at is task loss, the tasks changing."*

Marc Andreessen, Lenny's Podcast, 29 January 2026. His own caveat, minutes later: "as the tasks change enough, then that's when the jobs change." [53]

The idea comes from labour economics, which studies jobs as bundles of tasks. How you count makes a big difference:

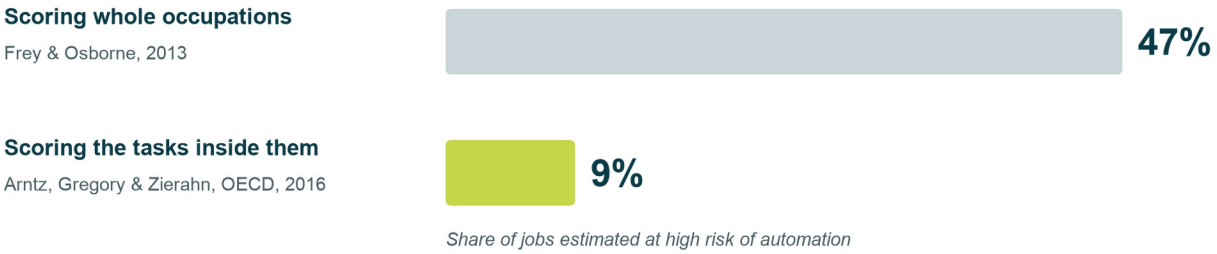


Figure 15 How you count changes the answer: the same question, scored by occupation and by task. [54, 55]

So far there is no sign of economy-wide job loss: Yale's Budget Lab finds no clear disruption through August 2026. But the first effects show up where you would expect. [56]

19%
below trend: US employment of 22 to 25 year olds in AI-exposed occupations [57]

Hiring
is where it shows up: fewer hires rather than layoffs, mainly where AI automates tasks rather than assisting

A map for every task

Anand and Wu (HBR, 2025) offer the most useful tool for deciding where AI goes. Break each job into tasks, then place each task on two questions: how costly would a mistake be, and does the task run on explicit data or on tacit, experience-based knowledge? [15]

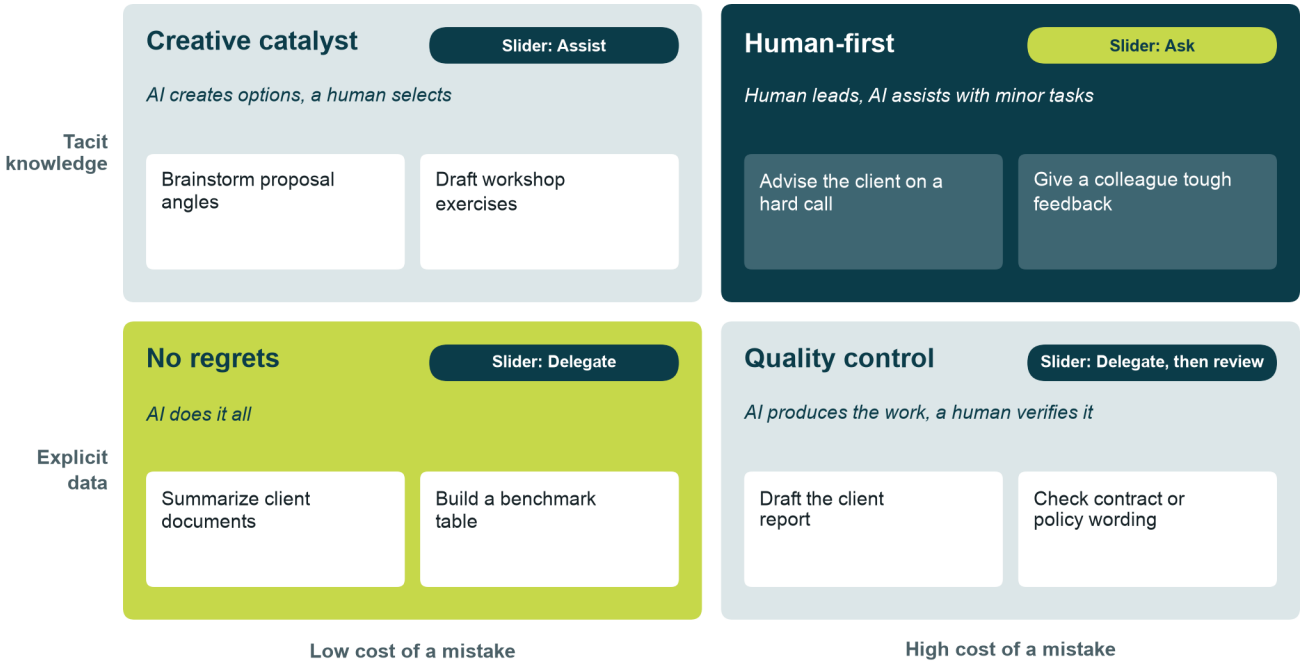


Figure 16 The task map, after Anand and Wu. The framework is theirs; the example tasks and slider settings are ours.

Many firms already decompose jobs into tasks for their clients; Ravin Jesuthasan and John Boudreau set out the method in Reinventing Jobs (2018). What has changed is which box each task lands in, and how quickly tasks move between boxes as models improve. [58]

The pyramid, pricing and competition

If AI takes a large share of an analyst's tasks, firms do not hire for fewer tasks. They hire fewer analysts.

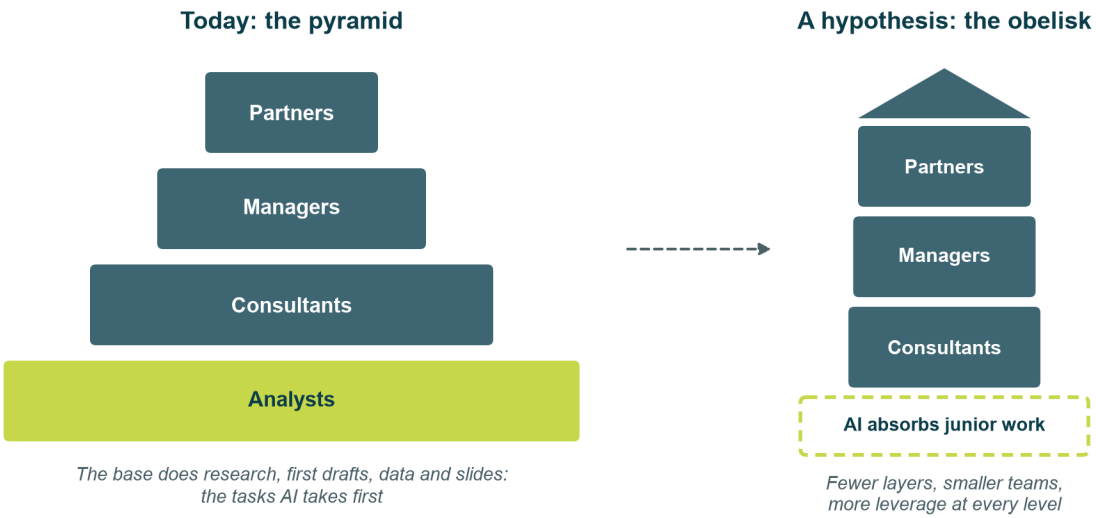


Figure 17 Some HBR authors expect the consulting pyramid to become an “obelisk”. Treat it as a hypothesis, not a forecast. [59]

Three questions follow for any firm built on leverage:

Who learns?

Juniors learn the craft by doing exactly the tasks AI takes first. If they become self-automators, they get fast without getting good.

Who pays?

When everyone has the tool, savings tend to flow to clients: Anand and Wu call it the paradox of access. McKinsey’s CEO said in January 2026 that about a third of the firm’s revenue is now tied to outcomes. [60]

Who competes?

Leaner firms, and AI labs that now sell implementation themselves. In May 2026 Anthropic, Blackstone, Hellman & Friedman and Goldman Sachs announced a new AI services company. [61]

Law shows the pricing pattern: US in-house counsel tripled from 1997 to 2020, and Anand and Wu report that nearly 90% of large law firms offer flat-fee or other client-favourable pricing, although such arrangements are still a small share of revenue at the largest firms. McKinsey’s CEO describes his firm as about 40,000 people and 20,000 agents.

> **The upside:** clients face the same redesign and will need help with it. Workforce redesign, AI governance, AI risk and the measurement of value are all services.

CHAPTER 8

Mental models to keep

Simple pictures that hold up under pressure.

The vehicle

"AI" covers very different things. Always ask which kind.

Narayanan & Kapoor, AI Snake Oil (2024)

Rulebook vs recipe

Software follows rules people write; AI learns patterns from examples, so it must be judged continuously.

Kozyrkov (2018); Karpathy, Software 2.0 (2017)

The brilliant new hire

Well read, fast, never met your client, rarely says "I don't know". Brief it, give it the file, check its work.

After Evans (2023) and Mollick (2024)

Cheap prediction, valuable judgment

AI lowers the cost of prediction. Deciding what to do with a prediction is still a human job.

Agrawal, Gans & Goldfarb (2017, 2022)

The autonomy slider

For each task, choose how much AI does alone: suggest, draft or act.

Karpathy (2025)

The jagged frontier

AI's strengths don't follow human difficulty. Test on your own work, and re-test often.

Dell'Acqua, Mollick et al. (2026)

The review tax

When drafting gets cheap, checking becomes the bottleneck.

OpenAI GDPval (2025)

The conversion gap

Time saved by people is not value captured by the firm until someone decides what the time is for.

NBER (2026); Ranganathan & Ye (2026)

New motor, old floor

Value comes from redesigning the work, not from bolting AI onto old processes.

David (1990); Hammer (1990)

The task map

Place each task by cost of error and type of knowledge before deciding AI's role.

Anand & Wu, HBR (2025)

Proof of value

Agree on the baseline and the outcome before anyone builds. A working demo is not a result.

Jaime Díaz-Beltrán

IF YOU KEEP ONE

Two redesigns beat twenty pilots

Pick two workflows, redesign them end to end, and agree on the outcome before anyone builds.

CHAPTER 9

A 90-day plan

Two redesigns beat twenty pilots. A simple plan for a practice or a firm.

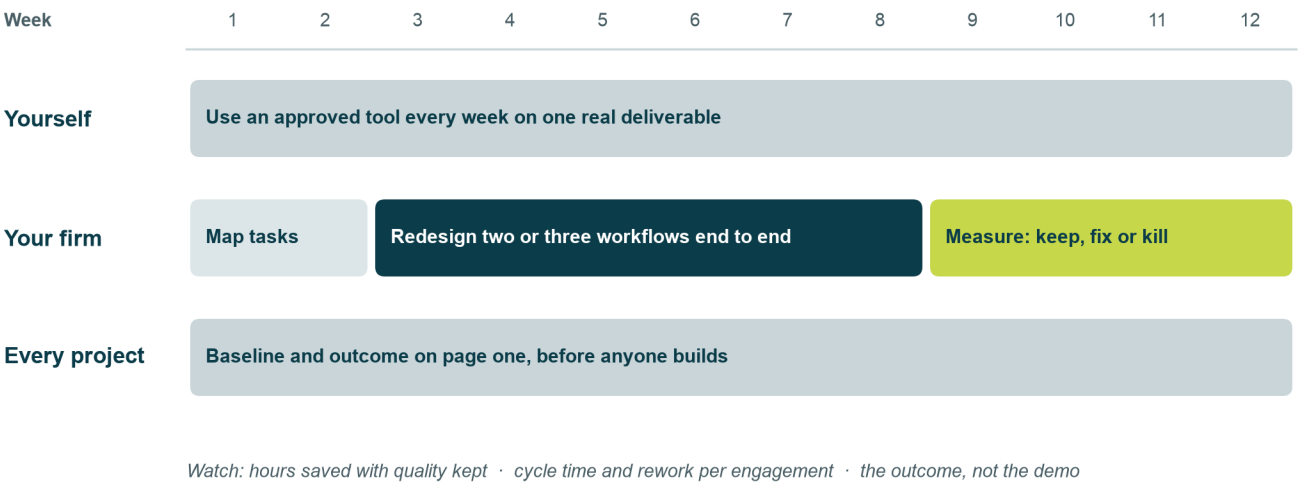


Figure 18 Ninety days, three moves, each with a number to watch.

1 · YOURSELF

Use it every week on real work

Use an approved AI tool every week on one real, recurring deliverable. Leaders who model use raise their teams' results. [62]

Watch: hours saved on that deliverable, with quality held

2 · YOUR FIRM

Redesign two or three workflows

Pick them with the task map. Redesign each one end to end in 90 days, instead of adding AI to the old steps.

Watch: cycle time and rework per engagement

3 · EVERY PROJECT

Agree on the outcome first

Write the baseline and the outcome on page one, before anyone builds. Proof of value, not proof of concept.

Watch: the outcome metric, not the demo

Week by week

- **Weeks 1 to 2:** map the tasks in two or three candidate workflows; agree on data rules and which tools are approved.
- **Weeks 3 to 8:** redesign the workflow around the no-regrets and quality-control tasks; decide who signs off at each step.
- **Weeks 9 to 12:** measure against the baseline; keep, fix or kill. Decide what the freed time is for.

CHAPTER 10

Where to go deeper

If you want to go further, start with five things that take about 18 hours in total.

Start here

What	Time	In Spanish
3Blue1Brown, "Large Language Models explained briefly" (video, 2024)	8 min	Spanish subtitles
Stanford HAI, "Inside the AI Index: 12 takeaways from the 2026 report"	15 min	English only
Ethan Mollick, Co-Intelligence (2024)	about 5 h	Cointeligencia (Conecta, 2025)
Arvind Narayanan & Sayash Kapoor, AI Snake Oil (2024)	about 7 h	No Spanish edition found
Andrew Ng, Generative AI for Everyone (DeepLearning.AI / Coursera)	about 6 h	Spanish dubbing on Coursera

Books

Book	Why read it	In Spanish
Agrawal, Gans & Goldfarb, Power and Prediction (HBR Press, 2022)	AI as cheap prediction; why system redesign beats point fixes	Prequel: Máquinas predictivas
Ethan Mollick, Co-Existence (Portfolio, Oct 2026)	The sequel to Co-Intelligence, for the agent era (new; no reviews yet)	Not yet
Andrew McAfee, The Geek Way (Little, Brown, 2023)	The culture side of adopting new technology	AI estilo geek
Melanie Mitchell, Artificial Intelligence: A Guide for Thinking Humans (2019)	The most honest plain account of what AI is and isn't	Inteligencia artificial: guía para seres pensantes
Parmy Olson, Supremacy (2024)	The OpenAI vs DeepMind race. FT & Schroders Business Book of the Year 2024	Not found
Stephen Witt, The Thinking Machine (2025)	Nvidia, chips and capital. FT & Schroders Business Book of the Year 2025	La máquina pensante
Karen Hao, Empire of AI (2025)	The cost side: labour, data, energy, power. Critical by design	El imperio de la IA
Mustafa Suleyman, The Coming Wave (2023)	An insider's case for containing AI's risks	La ola que viene

Courses

- Andrew Ng, AI For Everyone and Generative AI for Everyone (DeepLearning.AI on Coursera): non-technical, with Spanish dubbing.
- Anthropic, AI Fluency: Framework & Foundations (free, 3 to 4 hours): how to decide what to delegate to AI and how to check it. Made by a vendor, co-designed with academics.
- University of Helsinki, Elements of AI (free, Spanish version): the most rigorous course that needs no coding.
- Cassie Kozyrkov, Making Friends with Machine Learning (free on YouTube): how managers should frame AI projects as decisions.
- Executive programmes: MIT Sloan and CSAIL, “Artificial Intelligence: Implications for Business Strategy” (6 weeks online); Oxford Saïd, “Oxford Artificial Intelligence Programme”; Kellogg, “AI Strategies for Business Transformation”.

Videos

- The Thinking Game (Google DeepMind documentary, 1 h 24 min, free on YouTube, Latin American Spanish subtitles).
- Andrej Karpathy, “Intro to Large Language Models” (1 hour, 2023) and “How I use LLMs” (2025).
- Dot CSV, the leading Spanish-language AI explainer channel on YouTube.

To keep up

- Ethan Mollick’s newsletter, One Useful Thing.
- Stanford HAI’s annual AI Index, and Our World in Data’s AI section.
- METR and Epoch AI, for independent data on how fast capabilities and compute are growing.
- CEPAL and CENIA, Índice Latinoamericano de Inteligencia Artificial (ILIA 2026), for the regional view: Colombia ranks 4th of 20 countries.

REFERENCE

Sources

Numbered as cited in this document. Accessed between 28 and 29 October 2026 unless noted.

1. Weizenbaum, J. (1966). ELIZA: A Computer Program for the Study of Natural Language Communication Between Man and Machine. Communications of the ACM, 9(1), 36–45.

2. Lighthill, J. (1973). Artificial Intelligence: A General Survey. In Artificial Intelligence: A Paper Symposium. UK Science Research Council.

3. Krizhevsky, A., Sutskever, I. & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. NeurIPS. papers.nips.cc/paper_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html

4. Vaswani, A. et al. (2017). Attention Is All You Need. arxiv.org/abs/1706.03762

5. OpenAI (30 Nov 2022). Introducing ChatGPT. openai.com/index/chatgpt

6. Turing, A. M. (1950). Computing Machinery and Intelligence. Mind, 59(236), 433–460. doi:10.1093/mind/LIX.236.433

7. McCarthy, J., Minsky, M. L., Rochester, N. & Shannon, C. E. (31 Aug 1955). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. commons.wikimedia.org/wiki/File:A_Proposal_for_the_Dartmouth_Summer_Research_Project_on_Artificial_Intelligence,_by_John_McCarthy_et_al,_1955.pdf

8. Crevier, D. (1993). AI: The Tumultuous History of the Search for Artificial Intelligence. Basic Books.

9. IBM. Deep Blue. IBM Heritage. ibm.com/history/deep-blue

10. IBM. Watson, Jeopardy! champion. IBM Heritage. ibm.com/history/watson-jeopardy

11. Hassabis, D. (2016). What we learned in Seoul with AlphaGo. blog.google/technology/ai/what-we-learned-in-seoul-with-alphago

12. Jumper, J. et al. (2021). Highly accurate protein structure prediction with AlphaFold. Nature, 596, 583–589. doi:10.1038/s41586-021-03819-2

13. Nobel Prize Outreach (2024). The Nobel Prize in Physics 2024 and The Nobel Prize in Chemistry 2024, press releases. nobelprize.org

14. Google DeepMind (21 Jul 2025). Advanced version of Gemini with Deep Think officially achieves gold-medal standard at the International Mathematical Olympiad. deepmind.google

15. Anand, B. N. & Wu, A. (2025). The Gen AI Playbook for Organizations. Harvard Business Review, November–December 2025.

16. Hu, K. (2 Feb 2023). ChatGPT sets record for fastest-growing user base, analyst note. Reuters. reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01

17. OpenAI DevDay keynote (29 Sep 2026), as reported live by Engadget and TechCrunch.

18. Butters, J. FactSet Earnings Insight (26 May 2023; 14 Mar 2025; 12 Mar 2026; 18 Sep 2026): S&P 500 companies citing “AI” on earnings calls. insight.factset.com

19. Collins Dictionary (1 Nov 2023), Word of the Year 2023: AI. collinsdictionary.com/woty; Cambridge Dictionary (15 Nov 2023), Word of the Year 2023: hallucinate.

20. Tesler, L. Adages & Coinages. nomodes.com/Larry_Tesler_Consulting/Adages_and_Coinages.html

21. Narayanan, A. & Kapoor, S. (2024). AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference. Princeton University Press.

22. Kozyrkov, C. (2018). The simplest explanation of machine learning you'll ever read. Hackernoon. hackernoon.com/the-simplest-explanation-of-machine-learning-youll-ever-read-bebc0700047c

23. Karpathy, A. (11 Nov 2017). Software 2.0. karpathy.medium.com/software-2-0-a64152b37c35

24. 3Blue1Brown (2024). Large Language Models explained briefly. 3blue1brown.com/lessons/mini-llm

25. Karpathy, A. (Nov 2023). [1hr Talk] Intro to Large Language Models. youtube.com/watch?v=zjkBMFhNj_g

26. Google for Developers. Machine Learning Glossary. developers.google.com/machine-learning/glossary

27. Anthropic. Glossary. Claude documentation, docs.anthropic.com

28. NIST (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1). nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

29. OECD (2023, revised). Recommendation of the Council on Artificial Intelligence, definition of an AI system. legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

30. Karpathy, A. (17 Jun 2025). Software Is Changing (Again). Y Combinator AI Startup School. ycombinator.com/library/MW-andrej-karpathy-software-is-changing-again

31. OpenAI (2 Feb 2025) Introducing deep research and (17 Jul 2025) Introducing ChatGPT agent; Microsoft (21 Sep 2023) Microsoft 365 Copilot availability; Anthropic, Claude Cowork; OpenClaw (github.com/openclaw); Meta Muse (AP, 8 Sep 2026); OpenAI Dots (AP, 29 Sep 2026).

32. Russell, S. & Norvig, P. (1995). Artificial Intelligence: A Modern Approach. Prentice Hall.

33. Schluntz, E. & Zhang, B. (19 Dec 2024). Building effective agents. Anthropic. anthropic.com/engineering/building-effective-agents

34. METR (2026). Task-Completion Time Horizons of Frontier AI Models (Time Horizon 1.1). metr.org/time-horizons

35. Kwa, T., West, B. et al. (2025). Measuring AI Ability to Complete Long Software Tasks. arxiv.org/abs/2503.14499

36. Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz, H., Kellogg, K. C., Rajendran, S., Krayer, L., Candelon, F. & Lakhani, K. R. (2026). Navigating the Jagged Technological Frontier. Organization Science. doi:10.1287/orsc.2025.21838

37. Stanford HAI (2026). Artificial Intelligence Index Report 2026. hai.stanford.edu/ai-index/2026-ai-index-report

38. Randazzo, S., Lifshitz, H., Kellogg, K. C., Dell'Acqua, F., Mollick, E., Candelon, F. & Lakhani, K. R. (2025). Cyborgs, Centaurs and Self-Automators. HBS Working Paper 26-036. aiinstitute.hbs.edu/the-three-ways-professionals-work-with-ai-which-one-are-you

39. OpenAI, Anthropic, Microsoft and Google: enterprise data-use and training policies (accessed Oct 2026). Microsoft & LinkedIn (8 May 2024). 2024 Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part. KPMG & University of Melbourne (Apr 2025). Trust, attitudes and use of AI: a global study.

40. Associated Press (7 Oct 2025). Deloitte to partially refund Australian government for report with apparent AI-generated errors. Refund amount (A$97,587 of A$440,000) as disclosed at Australian Senate estimates, Oct 2025.

41. Mata v. Avianca, Inc., No. 22-cv-1461 (S.D.N.Y. 22 Jun 2023) (Castel, J.), opinion and order on sanctions; Corte Constitucional de Colombia, Sentencia T-323 de 2024.

42. Regulation (EU) 2024/1689 (Artificial Intelligence Act), Annex III, points 4 and 5(c); Regulation (EU) 2026/1744 (Digital Omnibus on AI), in force 27 Jul 2026, moving Annex III obligations to 2 Dec 2027.

43. Ley 1581 de 2012 (personal data protection); Superintendencia de Industria y Comercio, Circular Externa 002 de 2024 (21 Aug 2024), Tratamiento de datos personales en sistemas de inteligencia artificial; CONPES 4144 (2025), Política Nacional de Inteligencia Artificial.

44. Yotzov, I., Barrero, J. M., Bloom, N., Davis, S. J. et al. (2026). Firm Data on AI. NBER Working Paper 34836 (revised Oct 2026). nber.org/papers/w34836

45. BCG (30 Sep 2026). AI Is Starting to Pay Off. Almost 50% of Companies Now Generate Value with It; and BCG (30 Sep 2025). The Widening AI Value Gap. bcg.com

46. Korst, J., Puntoni, S., Tambe, P. et al. (Oct 2025). Accountable Acceleration: Gen AI Fast-Tracks Into the Enterprise. Wharton Human-AI Research & GBK Collective. ai.wharton.upenn.edu

47. Humlum, A. & Vestergaard, E. (2025, rev. 2026). Still Waters, Rapid Currents: Large Language Models and the Labor Market. NBER Working Paper 33777. nber.org/papers/w33777

48. Patwardhan, T. et al. (2025). GDPval: Evaluating AI Model Performance on Real-World Economically Valuable Tasks. OpenAI. arxiv.org/abs/2510.04374

49. Becker, J., Rush, N., Barnes, E. & Rein, D. (2025). Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity. METR. arxiv.org/abs/2507.09089

50. David, P. A. (1990). The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. American Economic Review, 80(2), 355–361.

51. Hammer, M. (1990). Reengineering Work: Don't Automate, Obliterate. Harvard Business Review, July–August 1990 (reprint 90406).

52. Agrawal, A., Gans, J. & Goldfarb, A. (2017). What to Expect From Artificial Intelligence. MIT Sloan Management Review; and Power and Prediction (Harvard Business Review Press, 2022).

53. Andreessen, M. (29 Jan 2026). The real AI boom hasn't even started yet. Lenny's Podcast, at 39:28. lennysnewsletter.com/p/marc-andreessen-the-real-ai-boom

54. Frey, C. B. & Osborne, M. A. (2013, publ. 2017). The Future of Employment: How Susceptible Are Jobs to Computerisation? Oxford Martin School; Technological Forecasting and Social Change, 114.

55. Arntz, M., Gregory, T. & Zierahn, U. (2016). The Risk of Automation for Jobs in OECD Countries. OECD Social, Employment and Migration Working Papers No. 189.

56. Yale Budget Lab (16 Jul 2026, updated 15 Sep 2026). Tracking the Impact of AI on the Labor Market. budgetlab.yale.edu

57. Brynjolfsson, E., Chandar, B. & Chen, R. (2026). Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Stanford Digital Economy Lab (rev. Aug 2026). digitaleconomy.stanford.edu

58. Jesuthasan, R. & Boudreau, J. W. (2018). Reinventing Jobs: A 4-Step Approach for Applying Automation to Work. Harvard Business Review Press.

59. Duncan, D. S., Anderson, T. & Saviano, J. (10 Sep 2025). AI Is Changing the Structure of Consulting Firms. HBR.org. hbr.org/2025/09/ai-is-changing-the-structure-of-consulting-firms

60. Sternfels, B. (6 Jan 2026). Where McKinsey and Consulting Go From Here. HBR IdeaCast, episode 1060. hbr.org/podcast/2026/01/where-mckinsey-and-consulting-go-from-here

61. Anthropic (4 May 2026). Building a new enterprise AI services company with Blackstone, Hellman & Friedman, and Goldman Sachs. anthropic.com/news/enterprise-ai-services-company

62. Microsoft (5 May 2026). Work Trend Index 2026.

Image credits

- Deep Blue at the Computer History Museum: Erik Pitti, CC BY 2.0, via Wikimedia Commons. Cropped and toned.
- Lee Sedol, May 2016: LG Electronics, CC BY 2.0, via Wikimedia Commons. Cropped and toned.
- AlphaFold predicted vs experimental structures: Jumper et al., Nature 596 (2021), Fig. 1, CC BY 4.0. Cropped and toned.
- Geoffrey Hinton and Demis Hassabis, Nobel week 2024: John Sears, CC BY-SA 4.0, via Wikimedia Commons. Cropped and toned; the edited images remain CC BY-SA 4.0.
- All charts and diagrams: drawn for this document from the sources cited in each caption.

> **A note on method.** Each claim was checked against its primary source where one was available (the paper, the official page or dataset, the court record), and then re-checked by a separate adversarial review. Quotes are verbatim unless marked as paraphrase. Where a figure rests on a single source or on a company with an interest in the result, the text says so.