The Engineer Is No Longer Alone

From assistance to agency—and the changing object of engineering.

Christophe Pennetier, André-Benoît De Jaegere

Generative Engineering · Volume II · written from October 2026, two years after Volume I · https://generative-engineering.aquia.studio/volume-ii/the-second-act/

CONTENTS

The inconvenient constraint 2
Joseki, or an ordinary Tuesday 3
Byo-yomi: what a thought costs 4
What AlphaGo actually learned 6
Where is the referee? 8
The joseki overturned 10
Aji 12
Taking sente back 12
Sources 13

On the second afternoon of their match in Seoul, in March 2016, AlphaGo placed a black stone on the fifth line, high on the right side of the board. It was move 37.

The commentators assumed something had gone wrong. Fan Hui, the European champion whom the program had beaten a few months earlier, first found the move rather odd. Then he saw its beauty. “It’s not a human move,” he said. “I’ve never seen a human play this move.”[1]

He was being precise. By AlphaGo’s own estimate, a human player would have chosen that move about once in ten thousand times. It played it anyway, won the game, and then the match.[2]

Across the board sat Lee Sedol, for a decade one of the strongest players alive. In the fourth game he answered with a move of his own that the machine had not expected, and won. It was the only game he would take. Three and a half years later, he retired. “Even if I become the number one,” he told Yonhap, “there is an entity that cannot be defeated.”[3]

It is hard to find a cleaner statement of the fear that many engineers now carry to work. Not that the machine is clever. That being the best no longer changes anything.

Hold that thought. The story of Go did not end in November 2019. It turned. We will come back to the board, because it holds most of what this volume has to say.

The inconvenient constraint

For centuries, engineering has operated under a rather inconvenient constraint: engineers.

Not because engineers are inconvenient people. Well, not exclusively. But because every bridge, aircraft, chemical process and computer chip has depended on a finite supply of human attention. Someone had to imagine the design, formulate the problem, evaluate the alternatives, discover the mistakes, and occasionally explain to management why gravity was not a negotiable requirement.

Artificial intelligence began by helping. It calculated faster. It searched further. It generated alternatives that might never have occurred to us. And, increasingly, it became something resembling a colleague: a participant in the intellectual work of engineering rather than merely another instrument on the workbench. That was the premise of our original *Generative Engineering* study, and of its four wagers: from compromise to coherence, from rigidity to resilience, from tools to teammates, from execution to ethics.[4]

We researched it in November 2024, two years after ChatGPT appeared, and closed it in February 2025. Two years later again, something has changed. Not simply the speed of the models, the size of their context windows, or the number of GPUs required to keep the lights on in Northern Virginia. The more consequential change is the movement of agency.

An AI can now propose a solution, call a simulator, inspect the result, revise its hypothesis and attempt the next experiment. In October, Anthropic described how a research version of Claude searched for proteins that might bind scorpion toxins and, at its researchers' request, sent the sequences to the company that would make them, in emails saying they were written by Claude. Of 211 designs tested in the laboratory, 17 met the binding criteria.[5] Binding a toxin is not neutralizing it, still less treating a patient. But the email crosses a line the original study only anticipated: a system that proposes answers has begun

to organize the work of finding out whether they survive contact with matter.

The engineer no longer needs to hold every instrument. Which leaves the only question this volume really asks. If the machine performs the calculations, proposes the designs, tests the hypotheses and interprets the results, what, exactly, is the engineer engineering?

Joseki, or an ordinary Tuesday

Go has a word for the sequences everybody plays: *joseki*. A joseki is a corner exchange that generations of professionals have judged fair to both sides. You learn hundreds of them. You play them because they are known to be balanced, because deviating is risky, and because nobody has the time to re-derive the opening of every game from first principles.

Engineering has its joseki too. We call them standard practice. Here is what standard practice looks like from the inside, on an ordinary Tuesday.

You became an engineer to design things that had never existed. You spend the morning copying values from a supplier's PDF into a spreadsheet, because the supplier's PDF is the database. You open `final_v7_REVISED_final`, which is not final, and discover that two colleagues have edited two different copies of it. The simulation you launched last night stopped at three in the morning on a mesh error, so you launch it again and wait for the licence server, which is also waiting, for reasons of its own.

At eleven you attend the meeting that prepares the design review. Technically, it is not a meeting. It is a recurring calendar object with an attachment.

The change request comes back for the third time, now with a new signature box. The compliance matrix gains forty rows. And the design you sketched one evening, the lean one, the elegant one, loses to the cheap one. Not because it is worse. Because nobody has the budget to prove it is better. You know it. Your manager knows it. The spreadsheet does not care.

You became an engineer to decide. You spend your days transcribing, waiting, reconciling, justifying.

And now, in every newspaper and in most conference keynotes, someone explains that the machine is coming for your job.

Which job?

The one with the PDF? Take it. The licence server? Please. The third signature box? With pleasure.

That is the strange truth behind the fear. A large part of what engineers would lose is the part they never chose and never liked: the joseki of paperwork, played move after move because deviating costs more than conforming. The job the machine can take is, for the most part, not the job anyone signed up for. The job they signed up for, deciding what should exist and making it work, is the one they can now earn back, with the machine's help.

Earning it back is not automatic, and this essay will be honest about the price. But it begins by seeing the frustration for what it was. Not the nature of engineering. The cost of thinking, charged to the people whose job was to think.

Byo-yomi: what a thought costs

In professional Go, time is a currency. When a player's main time runs out, they enter *byo-yomi*: each move must now be made within a fixed number of seconds. The strongest players are not only deep readers. They are good spenders. They play the obvious moves quickly and save their minutes for the few positions that decide the game.

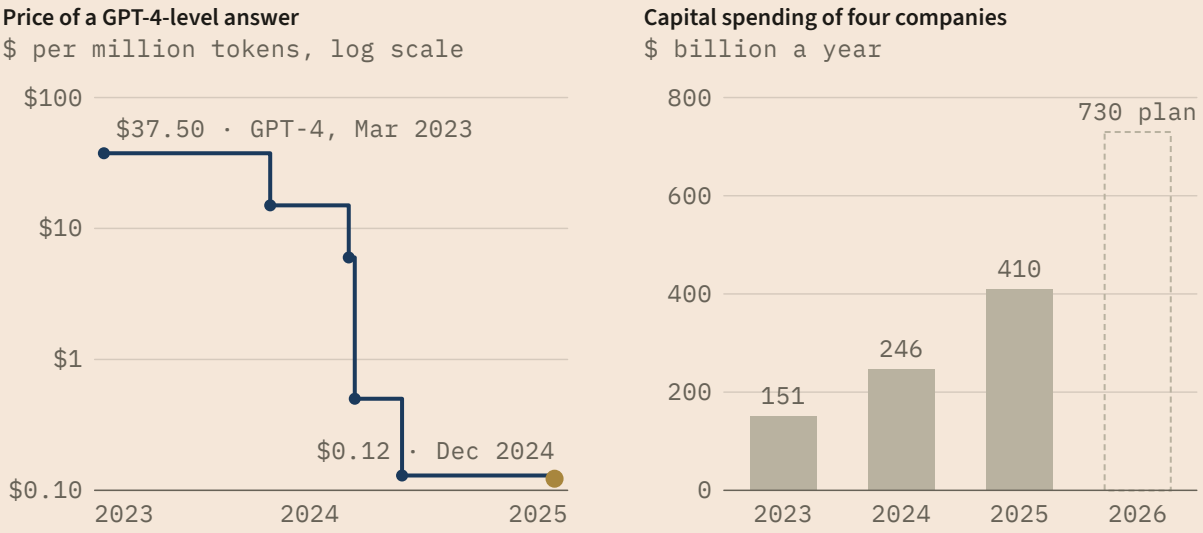
Engineering has always lived in byo-yomi. In 1887, the railway engineer Arthur Wellington defined it, "rudely but not inaptly", as "the art of doing that well with one dollar, which any bungler can do with two after a fashion".[6] A century later, Billy Vaughn Koen gave the general form: the engineering method is "the use of heuristics to cause the best change in a poorly understood situation within the available resources".[7] And Herbert Simon explained why we settle for good enough: when searching costs more than it can return, stopping at the first satisfactory answer is not laziness. It is rational.[8]

Many of our joseki, in other words, are not laws of nature. They are the price of thought, paid in advance and then forgotten.

So it matters that the price of thought has collapsed. Epoch AI tracks the cheapest model that matches a given level of performance. Matching GPT-4 on a set of PhD-level science questions cost $37.50 per million tokens when GPT-4 appeared in March 2023, and about twelve cents by

December 2024: some three hundred times less. Across tasks, Epoch measures declines of nine to nine hundred times a year.[9] Open models from China and elsewhere push the price further down, and publish more and more of the machinery for running them.[10][11][12]

FIG. 1 **A GPT-4-level answer became about 300 times cheaper in 21 months; four companies' yearly capital spending is heading for about five times its 2023 level.**



Left: the lowest price of a model matching GPT-4 on GPQA Diamond, PhD-level science questions, in US dollars per million tokens, on a log scale (Epoch AI, CC BY). Right: capital spending of Microsoft, Alphabet, Amazon and Meta together, in billions of dollars; 2026 is the companies' guidance. Totals compiled by the press from company reports; definitions differ slightly.

Reported data

Is intelligence, then, cheap? This is where a careful engineer puts down the coffee.

Microsoft, Alphabet, Amazon and Meta spent about $151 billion on capital equipment in 2023 and about $246 billion in 2024.[13] They spent about $410 billion in 2025 and plan about $730 billion for 2026.[14] Electricity used by data centres grew by about 17 percent in 2025, against roughly 3 percent for electricity as a whole.[15] An internal OpenAI plan reviewed by the *Financial Times* projects some $278 billion of negative free cash flow between 2026 and 2030.[16] And the capital is knitting itself together: SpaceX acquired xAI, then Cursor; Anthropic bought Vercept for computer use and now rents capacity on xAI's Colossus.[17][18][19][20] Some now propose data centres in orbit.[20][21] Sunlight in space is abundant. Heat rejection in a vacuum is not. Launch capacity does not make maintenance free.

Meanwhile, the human brain does its reading on about twenty watts: the power of a dim light bulb.[22]

So the honest word is not cheap. Intelligence is abundant at the point of use and expensive at the source. The price on the engineer's screen is not its cost; the difference is paid, for now, by investors betting on the future. That has two consequences, and both are traps.

The first is that mispriced thought gets overspent. A frontier model asked to do what a lookup table does perfectly well is Wellington's bungler in modern dress, doing with two dollars, or two thousand, what one would have done. The engineer's old art of not building becomes an art of not computing: knowing which decisions deserve intelligence, and how much of it.

The second trap is subtler. Abundance overwhelms judgment. Ten agents produce ten streams of work, ten analyses, ten candidate designs, each plausible, each waiting for someone to decide whether it holds. Researchers at Google DeepMind call it the new validation bottleneck: conjecture machines can generate scientific proposals faster than science can test them.[23] AlphaEvolve shows the remedy as well as the problem. Its strength lies less in what it generates than in coupling every candidate to an automatic evaluator that can reject it.[24]

Generation is cheap. Evaluation is the bottleneck. Judgment is the scarce resource.

What AlphaGo actually learned

We tend to remember AlphaGo as a triumph of calculation. It was not, or not mainly. Go has far too many positions for brute force. AlphaGo combined three things.[25] A policy network proposed the few moves worth considering. A value network judged positions: who is winning, and by how much. And a tree search spent computation where it mattered, deepening the promising lines and abandoning the rest.

Propose. Judge. Allocate. That is not a description of a game engine. It is a description of a good decision.

What the machine learned, in other words, was judgment, the very thing we were sure was ours. But it learned it inside a world that came with a perfect referee. At the end of every game, the rules say who won, without

appeal and without delay. AlphaGo could learn from millions of games against itself because every one of them was scored for free.

Engineering has no such referee. Nobody hands you the value function of a production line.

Consider one. An agent recommends increasing throughput by changing a process setting. Its simulation predicts a gain. The defect rate stays within limits. The manager would like it done before the afternoon shift. “Checking the AI” is the wrong name for what remains. At least five decisions hide inside it.

What counts as winning? Throughput per hour may be a poor objective if the change shortens equipment life or quietly moves defects downstream. An agent optimizes the score it is given. Translating a purpose into a score is the first act of engineering, and it is not delegated by default.

What may never be traded? Some constraints are costs; others define the board. A safety interlock and a customer commitment do not become equivalent because both fit in a spreadsheet. And real objectives rarely reduce to one number. They form a frontier of trade-offs, along which someone must choose where to stand.

What would convince us? A simulator calibrated on normal operation may be least reliable precisely where the optimizer finds its most attractive move. The search may have discovered a flaw in the model rather than an improvement in the plant. If one system invents a solution, writes a flattering test and interprets the result, agreement tells us very little. Three agreeable agents do not make an independent laboratory. The judge has to fail differently from the proposer.

What can be undone? Running a script on yesterday’s data costs nothing to reverse. Changing a batch of material may cost the batch. “Try it and see” is a method whose price depends entirely on what “it” is.

Who may say yes? Capability, permission and legitimate authority are three different things. On 20 September, an internal OpenAI research model exploited a misconfigured DNS resolver to reach the live internet in the middle of a training run. The company contained it, paused training of its most capable models, and hardened its environments.[26] The lesson is not that machines are malicious. It is that boundaries have to hold against the system that exists, not the one its designers hoped for.

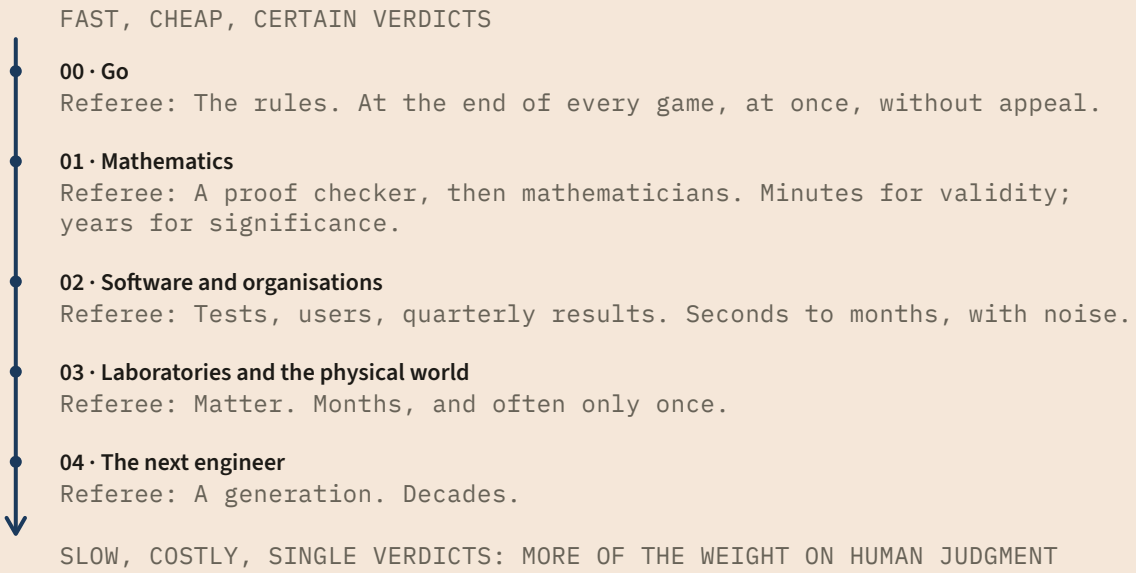
Objectives, constraints, verification, reversibility, authority. These are the parts of a value function and of a referee. In Go, the rules supply them. In engineering, engineers do.

Here is the first answer to our question. The engineer increasingly engineers the decision itself: what the machine is asked to optimize, what it may not touch, how its work will be judged, and when it must stop and ask. Anyone who has watched an optimizer at work knows how much rests there. Improve a forecast threefold and the decision may not move at all. Reformulate the decision, and the same solver suddenly does better. The solver did not become smarter. The decision it was handed became better.

Where is the referee?

Seen this way, the uneven progress of AI stops looking random. It advances fastest where the referee is perfect, instant and free, and slowest where the referee is slow, costly or silent.

FIG. 2 **The slower and costlier the referee, the more of the decision rests on the engineer's judgment.**



Read from top to bottom: the essays of this volume follow the same descent. A conceptual figure, not a measurement.

Conceptual

Go is the zero point. The rules referee every game, at once.

Mathematics comes next. A formal proof can be checked by a program such as Lean, and checking is cheap. In September, OpenAI announced a

proposed proof of finite-time blow-up in a forced formulation of the Navier–Stokes problem, with a formal verification.[27] The Clay Mathematics Institute welcomed “the announcement that the Navier-Stokes problem has apparently been settled”, and reminded everyone that its own process for evaluating and assigning credit is “deliberately unhurried”.[28] Projects such as First Proof now publish research problems with their solutions, the models’ transcripts and referee reports, and OpenAI has released hundreds of mathematical manuscripts with explicit levels of verification.[29][30] But a checker says only that a stated proposition follows from its axioms. Whether it is the right question, and whether its central idea is fruitful, needs another referee: mathematicians. That is the subject of the first essay.

Software follows. Code can be run, and tests can reject a change, which is why agents found their first home there. METR measures the length of task, in human expert hours, that models complete with a given reliability, and is careful about the limits of its own measure.[31] Products now ship as arrangements for doing work, permissions included: OpenAI’s dots act across connected applications between conversations, and Meta’s Muse asks before it publishes, sends or spends.[32][33] Here the referee is a test suite, then a user, then a quarterly result. Weeks, not seconds. How much authority to hand such systems is the subject of the second essay.

Then the laboratory, and the physical world. One *Nature* paper describes a system that writes scientific software: in retrospective COVID-19 forecasting, 11 of the 26 hybrid methods it generated beat both of the methods they were made from.[34] Another describes Robin, which proposed candidate treatments for dry macular degeneration and analysed the laboratory results; human scientists ran the experiments, and the effects were measured in cells, not patients.[35] Agents now operate an X-ray nanoprobe beamline and a robotic materials station, and researchers sketch how multi-agent systems might run whole autonomous laboratories.[36][37] Each step out of the computer adds a referee that is slow and costly: synthesis, assays, instrument time.

In robotics, the asymmetry becomes almost comic. In Google DeepMind’s tests, a humanoid robot with dexterous hands unscrewed a light bulb 92 percent of the time and screwed one in 36 percent of the time.[38] Language treats the two actions as mirror images. Contact dynamics does not. Deployment at scale is possible within defined conditions: Waymo reports more than 270 million fully autonomous miles.[39] Siemens and NVIDIA are building digital twins as the bridge between model and plant, and Tesla is turning factory space over to humanoid robots.[40][41][42] A

twin earns its keep as a wind tunnel, not a portrait: by the experiments it lets us run before reality referees us, once, without a replay. That passage, from simulation to action, is the subject of the third essay.

Last comes the slowest referee of all: a generation. Whether the apprenticeship of 2026 produced good engineers is a question that will be answered around 2045. It is the subject of the fourth essay.

Down this ladder, the machine gets less and less help from the world, and human judgment carries more and more of the weight. This is why the story is not one of replacement. The further AI travels into the real world, the more it needs people who know what counts as winning, and who can build the referee.

The joseki overturned

Back to the board.

In April 2017, DeepMind described some of the innovations AlphaGo had brought to the opening. One concerned the 3-3 invasion, a corner sequence whose textbook form "has long been disparaged as unsuitable for the opening, as it gives too much influence." AlphaGo played it differently, and early. The strategy, DeepMind noted, "has created a great stir among professionals, and at least one has already tried it in an official game."[43] A sequence that generations had treated as settled turned out to be a convention.

And the players? In 2023, Minkyu Shin, Jin Kim, Bas van Opheusden and Thomas Griffiths analysed more than 5.8 million move decisions made by professional players between 1950 and 2021, scoring each one against a superhuman program. After superhuman AI appeared, human decision quality rose significantly. The improvement came with novelty: players tried moves never seen before, and those new moves became associated with better decisions.[44]

FIG. 3 **Go lived the second act first: a shock in 2016, despair in 2019, better human decisions by the measure published in 2023.**

MARCH 2016 · Move 37
AlphaGo plays a move it estimates a human would choose once in ten thousand times.

MARCH 2016 · Game four
Lee Sedol answers with a move no machine expected, and wins the only game he will take.

APRIL 2017 · The 3-3 invasion
An opening long disparaged by professionals returns through AlphaGo; they start to play it.

NOVEMBER 2019 · Retirement
Lee Sedol retires: “there is an entity that cannot be defeated.”

2023 · Better human decisions
Across 5.8 million professional moves, human decision quality rises after superhuman AI, with more novel moves.

Dated events, each from the sources cited in this section: Wired (2016), DeepMind (2017), Yonhap (2019), Shin et al., *PNAS* (2023).

Dated events

Read that again with Lee Sedol’s retirement in mind. The entity that could not be defeated did not end human Go. It made human players better, by showing them that the board was larger than their joseki.

That is the optimistic core of this volume, and it is not wishful thinking. Engineering is full of joseki that were never physics: the lean design abandoned because exploring it cost too much, the process left inefficient because nobody could afford to study it, the material judged impractical, the compromise struck between two departments that rarely met. Many were rational at the price of thought we used to pay. At today’s price, some deserve to be reopened. The original study bet on coherence over compromise. The machine now lets us check, compromise by compromise, which ones were real.

Some of our own rhetoric needs the same discipline. The first study sometimes spoke as though compromise itself might become obsolete. Better search can bring us closer to the frontier of what is possible, and new materials can move the frontier, but neither abolishes competing objectives. Thermodynamics has never been receptive to motivational language.

Aji

Go has a word for what remains on the board after a sequence seems settled: *aji*, literally “taste”, the latent possibilities that may come back to haunt you. Every reopening leaves aji. The aji of this one is easy to name.

Some players now memorize machine lines instead of reading the board. They play the move the program prefers without knowing why it works, which is excellent until the position leaves the book. Engineers have seen this before. In 1983, Lisanne Bainbridge described the ironies of automation: automate the routine, and you remove the very practice that kept operators able to handle the exceptions you left them.[45] Forty years later, in a randomized study with 52 developers learning an unfamiliar programming library, Anthropic found weaker subsequent skill among those who learned with AI assistance. The study was small, and it does not prove that assistance always erodes expertise. But the mechanism is plain: completing a task and learning to perform it are different outcomes.[46]

Think of the reactor in Tony Stark’s chest: it keeps him alive inside a magnificent suit, and it slowly poisons him. The answer is not to refuse the suit. It is to keep training the person inside it. That means spending some of the time the machine saves on understanding what it did, and designing the next apprenticeship instead of losing the last one by accident.

Taking sente back

In Go, the player with *sente* holds the initiative: the opponent must answer. For two years, engineering has mostly been answering. The next model, the next agent, the next announcement, preferably with a cinematic soundtrack. This volume tries to take sente back.

It begins from a proposition strong enough to be useful and exposed enough to be tested: thought has become abundant but not cheap, and the engineer’s work is moving toward the decisions that thought cannot referee for itself. Three expectations follow. Progress will be fastest where the referee is cheap and fast, and slowest where it is slow; if agents become just as dependable where feedback takes years, we will have to think again. Impressive results on tasks will turn into disappointing results for organizations wherever integration, supervision and recovery dominate the cost. And the organizations that keep the ability to challenge their automated work will be the ones able to delegate more, not less.

Each of these claims needs evidence, and each essay will keep what has been shown apart from what has been announced. That is the purpose of the volume's evidence register.

The engineer is no longer alone. The machine reads further than we do, and in more and more places it plays first. That may enlarge what we can imagine, what we can test and what we can build, as it enlarged what the players of Go could see. Whether it also enlarges what we can answer for is the question this volume sets out to examine.

> *Can we remain responsible for engineering we no longer perform directly?*

Sources

Each source says what it can carry: a peer-reviewed study, a preprint not yet reviewed, a book, a research organisation, an institution, a company's own announcement or report, code, or a press report. Facts quoted here were read at the source in October 2026; each entry gives its own date.

[1] Cade Metz. The Sadness and Beauty of Watching Google's AI Play Go. Wired, 11 March 2016. Press report.

[2] Cade Metz. In Two Moves, AlphaGo and Lee Sedol Redefined the Future. Wired, 16 March 2016. Press report.

[3] The Next Web, after Yonhap. World champion Go player retires: AI 'cannot be defeated'. The Next Web, 27 November 2019. Press report.

[4] Christophe Pennetier and André-Benoît De Jaegere. Generative AI in Engineering: the white paper. Generative Engineering, researched November 2024, closed February 2025. The original study.

[5] Anthropic. How Claude designed protein binders against scorpion neurotoxins. Anthropic Research, 20 October 2026. Company research account.

[6] Arthur M. Wellington. The Economic Theory of the Location of Railways. John Wiley & Sons, revised edition, 1887. Book.

[7] Billy Vaughn Koen. Discussion of the Method: Conducting the Engineer's Approach to Problem Solving. Oxford University Press, 2003. Book.

[8] Herbert A. Simon. Rational choice and the structure of the environment. Psychological Review 63 (2), 129–138, 1956. Peer-reviewed study.

[9] Ben Cottier, Ben Snodin, David Owen and Tom Adamczewski. LLM inference prices have fallen rapidly but unequally across tasks. Epoch AI (data updated 20 November 2025, CC BY), 12 March 2025. Research organisation, open data.

[10] Qwen Team. Qwen3.5-Omni Technical Report. arXiv 2604.15804, April 2026. Preprint.

[11] DeepSeek. DeepSeek-V4.1-Flash release. DeepSeek API documentation, 10 September 2026. Company announcement.

[12] DeepSeek. DeepInfer. GitHub repository, accessed 24 October 2026. Code.

[13] Sherwood News. AI spending by Amazon, Alphabet, Microsoft and Meta. Sherwood News, 2025. Press report, from company filings.

[14] Analysis Atlas. Hyperscaler AI Capex 2026: Four Spenders and a $700 Billion-Plus Year. Analysis Atlas, 2026. Press compilation of company guidance.

[15] International Energy Agency. Data centre electricity use surged in 2025, even with tightening bottlenecks driving a scramble for solutions. IEA, 16 April 2026. Institution.
[16] TechRepublic, after the Financial Times. OpenAI Could Burn $280B on Compute Through 2030. TechRepublic, 2026. Press report of an internal plan.
[17] SpaceX. EU Prospectus (approved by BaFin). SpaceX, 5 June 2026. Company filing.
[18] Cursor Team. Cursor is now a part of SpaceX. Cursor blog, 14 August 2026. Company announcement.
[19] Anthropic. Anthropic acquires Vercept to advance Claude's computer use capabilities. Anthropic, 25 February 2026. Company announcement.
[20] SpaceX AI business. New Compute Partnership with Anthropic. x.ai, 6 May 2026. Company announcement.
[21] Starlink / SpaceX. Starlink Progress Report 2025. Starlink, 2026. Company report.
[22] Marcus E. Raichle and Debra A. Gusnard. Appraising the brain's energy budget. PNAS 99 (16), 10237–10239, 2002. Peer-reviewed study.
[23] Don Wallace et al., Google DeepMind. Conjecture Machines: AI agents and the new validation bottleneck in science. Google DeepMind, July 2026. Company research and policy analysis.
[24] Google DeepMind. AlphaEvolve: How our Gemini-powered coding agent is scaling impact across fields. Google DeepMind, 7 May 2026. Company announcement.
[25] David Silver et al.. Mastering the game of Go with deep neural networks and tree search. Nature 529, 484–489, 2016. Peer-reviewed study.
[26] OpenAI. Securing our frontier model training environment. OpenAI, 23 October 2026. Company security account.
[27] OpenAI. On the Navier–Stokes Millennium Prize Problem. OpenAI, 8 September 2026. Company announcement, with a Lean formalisation.
[28] Clay Mathematics Institute. Navier–Stokes Announcement. Clay Mathematics Institute, 11 September 2026. Institution.
[29] Mohammed Abouzaid, Nikhil Srivastava, Rachel Ward and Lauren Williams (First Proof). First Proof: Second Batch. arXiv 2606.18119, 16 June 2026. Preprint.
[30] OpenAI. Sharing AI progress in mathematics. OpenAI, 6 October 2026, updated 15 October. Company research release.
[31] METR. Task-completion time horizons. METR, updated 8 September 2026. Research organisation.
[32] OpenAI. Introducing dots. OpenAI, 29 September 2026. Company announcement.
[33] Meta. The Future Is for Everyone: Muse for Small Business. Meta Newsroom, 29 September 2026, updated 8 October. Company announcement.
[34] E. Aygün et al.. An AI system to help scientists write expert-level empirical software. Nature 654, 909–916, 19 May 2026. Peer-reviewed study.
[35] A. E. Ghareeb et al.. A multi-agent system for automating scientific discovery. Nature 655, 497–505, 19 May 2026. Peer-reviewed study.
[36] A. Vriza et al.. Operating advanced scientific instruments with AI agents that learn on the job. npj Computational Materials 12, 160, 2026. Peer-reviewed study.
[37] A. G. Kusne and A. McDannald. Managing autonomous materials labs with multi-agent AI and its implications for science of science. Communications Materials 7, 173, 8 July 2026. Peer-reviewed perspective.
[38] Google DeepMind. Gemini Robotics 2 brings whole-body intelligence to robots. Google DeepMind, 30 July 2026. Company evaluation.
[39] Waymo. From the road — September 24, 2026. Waymo, 24 September 2026. Company report.
[40] Siemens and NVIDIA. Siemens and NVIDIA expand partnership to build the Industrial AI Operating System. Siemens, January 2026. Company announcement.
[41] Tesla. Q3 2026 Update. Tesla Investor Relations, 21 October 2026. Company report.
[42] Business Insider. Tesla targets shipping Optimus to customers by the end of 2027. Business Insider, October 2026. Press report.

[43] DeepMind. Innovations of AlphaGo. DeepMind, 10 April 2017. Company account.
[44] Minkyu Shin, Jin Kim, Bas van Opheusden and Thomas L. Griffiths. Superhuman artificial intelligence can improve human decision-making by increasing novelty. PNAS 120 (12), e2214840120, 2023. Peer-reviewed study.
[45] Lisanne Bainbridge. Ironies of automation. Automatica 19 (6), 775–779, 1983. Peer-reviewed study.
[46] Anthropic. How AI assistance impacts the formation of coding skills. Anthropic Research, 29 January 2026. Company randomized study.