What is needed to create a $1T AI-driven materials discovery company?

Sina Shahandeh

November 2026

If only we could ask Claude to design new materials. We would build foundation models for materials discovery, run them in parallel and try billions of candidates in thousands of processes and applications. Surely thousands of geniuses in a gigawatt data center [1] could solve many of humanity's problems and capture trillions of dollars of value along the way. Catalysts for any reaction, strong materials for space and transport, materials that store energy in batteries, superconductors and semiconductors, functional materials for everything from buildings to clothes to medical implants. There is no shortage of things people want.

Sadly, there is a catch. AI-driven materials discovery hits a bottleneck very early.

Candidates are easy

In 2023 Google DeepMind's GNoME predicted 2.2 million new crystal structures [2]. The same year an autonomous lab in Berkeley reported 41 new materials synthesized in 17 days [3]. In 2025 Microsoft's MatterGen designed a crystal to be hard to compress, aiming for a bulk modulus of 200 GPa [4]. Bulk modulus measures how hard a material is to squeeze, and steel's is about 160 [5].

On subsequent examination these claims became less stellar. Two solid-state chemists who went through a sample of the GNoME structures found scant evidence of compounds that were novel, credible and useful at the same time [6]. Crystallographers who re-examined the Berkeley lab's data argued that most of its new materials were known compounds in disordered form [7]. MatterGen's crystal scored 222 GPa on the computer. The sample that came out of the furnace had its tantalum and chromium atoms scrambled across sites the model had kept apart, and it measured about 160, close to ordinary steel. A hard material also has to do more than resist squeezing, and the MatterGen authors noted themselves that it must resist shear as well [4]. Each generator loop did the narrow job that it was designed to do. The trouble started afterwards, in the slow work of establishing what was real, what was new and what would work in the real world.

The lab in the loop

Lila Sciences and Periodic Labs saw this bottleneck early. They raised hundreds of millions of dollars to build what Lila calls "scientific superintelligence" [8], with a physical laboratory inside the loop of discovery [9, 10]. Robots make and test candidates at high throughput, the results flow back into the models, and the models design the next round with more signal and less noise. The idea is that, given enough rounds, the best candidates rise to the top.

The approach works. In September 2026 Lila reported screening 2,942 catalyst compositions for splitting water in acid, and finding palladium-oxide families that avoid iridium, the scarce metal today's acid electrolyzers depend on [11]. The best held its performance for a thousand hours, where earlier high-throughput screens had tested stability for thirty minutes [12]. A loop like this can plausibly make the design stage ten times more productive [13, 14].

However, design is only the first stage of a long journey. Lila's catalyst was tested in a small cell of sulfuric acid [11]. An industrial electrolyzer pushes 200 to 300 times more current through its catalyst and has to keep going for tens of thousands of hours [15]. Records set in small cells have often failed to survive scaled production [16, 17]. Fast experiments also measure one or two properties at a time. A material in production has to satisfy many at once, and they tend to pull in opposite directions: make a metal stronger and it usually becomes more brittle [18]. By nature, high-throughput experiments sacrifice the extent of what is measured for the speed and cost of the test. Ultimately, an agent in the loop can only optimize what it can measure.

The discovery ladder and the impact of AI

Drug discovery has been down this road. Candidate molecules are easy to generate, and the first tests, binding to a protein or killing cancer cells in a dish, are fast. The real test is the human body, where a molecule that binds to the target can fail to shrink the tumor inside the organ. About nine in ten drugs that enter human trials fail [19]. In the first phase of trials, which mostly tests safety, molecules discovered with AI now pass 80 to 90 percent of the time, against 40 to 65 percent historically. In the second phase, which asks whether the drug works, they pass about 40 percent of the time, no better than the historical average [20].

Materials have their own human body: the factory, and then years of service in the field. At the current rates a new material takes ten to twenty years to go from a promising sample to production by the ton [21, 22]. In the table below, we divide this journey into stages, with the clock starting once a candidate is perceived as promising.[1]

Down the table the share of the clock grows and the gains with AI shrink, as the limit moves from computation and ideas to physical work, then to a handful of expensive runs, then to institutions and product cycles.[2] In this dynamic, halving the scale-up stage saves almost twice as much time as one would save by making the lab go ten times faster.[3]

[1] The shares in the table are my estimates. The gains are what today's tools can plausibly deliver.

[2] Adoption's gain assumes that about half of its elapsed time is human-paced work that runs 1.5 to 3 times faster with AI, and that the rest is waiting on product cycles. Field experiments find speed-ups of 15 to 40 percent on tasks AI suits [35, 36], workers report saving 3 to 5 percent of their hours [37, 38], and executives report that AI has added 0.2 percent a year to labor productivity [39]. The low end of the table's range, 1.2×, is what today's evidence supports.

[3] At the table's midpoints the end-to-end factor is 1.84. Halving scale-up saves 12.5 points of the pipeline's time, and a tenfold faster lab saves 6.8. Appendix A gives the arithmetic.

Stage	Share of the 10–20 years	What makes it hard	What AI can do	Plausible gain	AI companies working here
Physical understanding	No clock of its own; it sets the odds at every stage	Simulations follow a few hundred atoms for trillionths of a second; cracking, corrosion and diffusion are out of reach [23, 24]	Learned simulators, orders of magnitude faster, still graded mostly against calculations [25, 26]	1,000× in what can be simulated [27]; unknown in accuracy	Google DeepMind, Meta, Microsoft [2, 4, 28]
Candidate generation	Under 5% where the physics is known; open-ended where it is not	A generator can aim only at what can be computed [29]	Generative design and stability filters within known chemistry [4]	Several-fold hit rates where the physics is known [26]	Most AI-for-materials startups
Lab synthesis and screening	5–10%	Throughput, and whether the test predicts anything	Robots, parallel synthesis and automated measurement [13, 14]; longer, more realistic tests [11]	10×	Lila, Periodic [8, 9]
Design validation	10–20%	Every property at once, in the real form and conditions, over months of testing	Predicting long tests from short ones [30, 31]; learning how cheap tests map to expensive ones	2–3×	Few
Scale-up	20–30%	A handful of pilot runs; know-how that lives in people	Models of the production line built from its own data, with physics as a guide	About 2×	Almost none
Qualification	25–35% in aerospace, less elsewhere	Customers and standards bodies ask for many full-scale batches	Models that stand in for some tests [32, 33]; qualifying a window of recipes at once	1.5×, more if standards bodies move	QuesTek, in steel [34]
Adoption	15–25%	Product cycles; requalifying a supply chain [22]	The productivity gain AI brings to any industry [35, 36]	1.2–1.5×	None specific to materials

Conceptually, this table shows a pipeline with a screening filter at each stage. How many initial candidates succeed depends on the quality of these screening tests along the way. In 2016 Jack Scannell and Jim Bosley showed, for drugs, that a screen's predictive validity, how closely its score tracks the real outcome, matters more than how many candidates it can process [40, 41]. It is the hiring problem: if the interview barely predicts performance, interviewing ten times more people mostly finds people who interview well. The difference is that in the materials (or drug) world the number of applicants is on the order of 10^{60}.[4] In a toy version with such numbers, where one candidate in a million works in the field and the lab score correlates 0.3 with field performance, a test ten percent more predictive is worth about as much as ten times more tests, and a test a third more predictive beats a hundred times more test bandwidth. With autonomous

[4]The figure is the usual estimate of the number of possible drug-like molecules with up to 30 atoms of carbon, nitrogen, oxygen and sulfur [42].

labs running thousands of samples a month, we are already in a regime where a further tenfold buys little. We need to improve the tests' screening ability. See Appendix A for the model math.

At the candidate generation stage, a good lab should also be able to be surprised. Stainless steel, Teflon and Kevlar were all noticed during other projects [43–45]. The iron catalyst behind synthetic fertilizer came out of some 6,500 runs over 2,500 formulations, when the impurities in one Swedish magnetite turned out to make it work far better [46]. Only an AI that reads every diffraction pattern, image and lab log could find surprises like these at scale. This differs from traditional automated labs, where the protocol measures only what it was asked to optimize.

Another important point is increasing the validity of the tests. At the lab synthesis and screening stage, a test becomes more predictive when the factory conditions are brought into the lab. A robotic lab holds a sample's temperature to within a fraction of a degree. A production furnace has hot and cold zones, and what fails at scale is often a material that could not tolerate them. The question is how we can build the furnace's temperature spread into the lab protocol, along with the impurities of industrial feedstock and the mixing of a large vessel, so that the lab test starts to predict the plant.

What is needed

Lab screening tests can become more realistic, yet there is a fundamental gap between materials design and production. In the spirit of a $1T company, let's look at an analogy. When NVIDIA's engineers finish a chip design, they already know TSMC can make it. Before a wafer is committed, the design is checked against rules and simulation models that TSMC calibrated on its own production line. So there is a filter that gives designers the certainty that what is drawn will be made functional. That certainty lets NVIDIA own the most valuable chips in the world without owning a factory, and lets TSMC build chips for other companies without designing any of its own. It took decades to build. Carver Mead and Lynn Conway published design rules of this kind in 1980, so that a designer could target a process they would never see [47]. Morris Chang founded TSMC in 1987 to make other people's designs while not competing with them [48]. Synopsys and Cadence grew up in between to sell the design tools, and they are now worth nearly $180 billion together [49].

Now picture a head of R&D at a specialty chemicals company with a coolant that beats the one the company sells. It has passed every test in the lab. To make it, the plant would have to stop a line that earns money every day, retune its reactors and run this new batch. Will it come out on spec? Will it foul the heat exchangers or clog the pipes? Nobody can say. The plant manager is paid to keep the line running. So the new coolant remains in a paper and the plant makes what it made last year.

The design rules split the risk of a new chip in two. The foundry's question, whether the chip can be made successfully, is answered before anything is built. The designer that keeps the other uncertainty, whether anyone will buy it, is paid for that risk with the product margin. This is why NVIDIA is worth more than twice as much as TSMC [50]. Currently, a materials maker carries both questions at once and is paralyzed over whether they should risk current production for an unproven product. In McKinsey's study of successful chemical launches, new technology for new markets took 8 to 19 years and returned 8 to 12 percent a year, about the cost

of capital, while known products in known markets returned 18 to 23 percent [51]. The CFO who says no to the head of R&D has done the arithmetic.

Materials today look like chips before TSMC. Corning, Carpenter, Umicore and Dow design their own materials and run their own plants. There are legends about how Steve Jobs convinced Corning's CEO to go all in on iPhone glass [52]. Startups with a better material, such as Sila and Group14 in battery anodes, have had to build plants of their own [53, 54]. Contract manufacturers, called tollers in the trade, will run a customer's recipe for a fee [55], but nobody can tell beforehand whether a new recipe will run on their line, so every transfer is a large capex experiment. The first fabless chip companies leaned on established makers in the same way. Xilinx's first chip, in 1985, came out of Seiko Epson's fab [56, 57].

Three things would give materials the certainty chips have.

The first is a physical understanding that is able to predict composition-process-microstructure-property links. The quantum calculations that computational materials science runs on can follow a few hundred atoms for a few trillionths of a second [23, 58]. Whether a material is useful depends on cracking, corrosion, diffusion and the slow drift of its microstructure, which play out over meters and decades [18, 24]. For some targets the physics is missing altogether. Forty years after the copper-oxide superconductors were discovered [59], nobody agrees why they work [60, 61]. Beyond perfect crystals, phenomena at interfaces and in the presence of impurities and multiple compounds remain unknown. It is astonishing simply how much we do not know.

The second is the understanding of the fabrication process. An example from my undergraduate years in metallurgy: on a forging line, the way lubricant is spread on a metal sheet under a press sets the friction. The friction sets the stress in the part, and the stress opens a crack when the part is in the furnace, hundreds of meters down the line. Cause and effect sit in different places, at different times, on different sensors. Only some links live in the heads of experienced people on the shop floor. Formalizing the fabrication process into an accurate model is the key here. Such a model can be built by accurate measurement and characterization of the production line, far beyond what is done today by traditional industrial IoT.

The third is a way to qualify a range. Qualification means showing a customer or a standards body, with many full-scale batches, that a material meets its specification. A new chip inherits the qualification of the process it is made on, while every new alloy or polymer is qualified on its own, one composition and one process at a time. If a line were understood well enough to qualify a window of compositions and processes, every design inside the window would inherit most of the evidence. Drug regulators already work this way: a manufacturer that demonstrates a "design space" can move within it without a new approval [62].

What this means for AI-for-science companies

Follow the money. Non-commodity materials sell about a trillion dollars a year. Their makers and governments spend about $57 billion a year on R&D, perhaps $9 to $17 billion of it on discovering new materials. Outside vendors sell $2 to $4 billion of R&D services into that budget, and materials-informatics software earns $45 to $115 million. About a billion dollars of venture

money a year now goes into AI for materials discovery.[5] R&D can benefit from AI, but the money in materials sits in production.

Ask the head of R&D at a chemical company what holds them back, and you will hear about a shelf of candidates that cannot get into the plant. Aionics began by selling electrolyte-design software to battery makers, found that customers could not put a price on faster research, and now develops formulations that partners manufacture [63].

So the place to start is the plant. Plants already record far more than they use. In 2015 two McKinsey researchers wrote that a modern oil platform carries over 30,000 sensors and that less than one percent of their data was used for decisions [64]. Add cameras and sensors where the gaps are, bring together the logs, maintenance records and operators' notes, and build a model of how the line really behaves: where the furnace runs hot, which valve drifts, which step turns a small variation into scrap. A few years ago this would have been too expensive. Today's AI models read video, images, spectra, sensor logs and handwritten notes together, which makes a model of a messy, sparsely observed line affordable.

For the head of operations, the model should pay for itself within months in yield, uptime and energy. It would be great if the contract gave the company that built the model the right to learn from the line. NVIDIA began in gaming [65], a market that paid for its graphics chips and CUDA long before AI became their largest use. For this AI company, participating in plant operations can be a beachhead that monetizes materials discovery.

The model that is built in this way has predictive power versus the one built in the lab alone: given this line, will the new material run on it, come out on spec and leave the reactors clean? With this model, the company can do what NVIDIA does. It can design a material, show the plant that it will run, pay for a month of production and sell the material under its own name. Suppose the material is a coolant better than anything on the market, in a niche worth half a billion dollars a year. The company would keep both the customer and the product margin.

The first markets should be those where the ladder from sample to production is short and the physics is well understood. Coatings, polymer compounds, catalysts for existing plants and the materials inside electronics qualify in about one to three years, while aerospace alloys take ten or more [22]. Data-center coolants may come first. Orbital Industries is taking its AI-designed coolant through qualification with major chip makers, aiming to ship it with the next generation of GPUs in 2027, which its chief executive says would make it the first AI-designed molecule on the market [66].

A $1T company is possible if. . .

As a rough estimate, new classes of materials outside batteries bring in $8 to $20 billion a year, about one to two percent of the trillion-dollar market for non-commodity materials. That share is small because a new material takes so long and costs so much to launch that it earns about the cost of capital [51]. If AI cut the time by the table's factor, a program that returned 10 percent

[5] All six figures are my estimates, built from company filings, national statistics and market reports. The ranges are $0.9–1.2 trillion for non-commodity sales and $41–79 billion for R&D, and venture funding was about $1.07 billion in 2025 with Lila Sciences included. Schrödinger, the one company that reports materials software revenue, reported $17 million for 2025.

would return 14 percent, or 22 percent if spending fell with the time.[6] If the plant were certain that a new material would run, the designer could carry the market risk alone. These two conditions would make new materials worth funding, and the market for them would grow.

So one $1T company is possible if four things happen. A model developed on paired lab and production data is able to predict, ahead of time, a durability result that takes a year to measure. A plant runs a material it has never made, at the designer's expense, and it comes out on spec the first time. A customer or a standards body accepts a qualified window of processing conditions. And a material shipped this way earns a product margin. Core pieces are here and each can be tested within three years. The company that owns the products and the line models could hold the seat NVIDIA holds in chips. The integrated modeling and foundry partners would be the plants it learned from.

Ten to twenty years from now I expect factories to be run fully by AI agents. The agent that runs a line would also be the same agent that runs its R&D, because the two are one loop. What the line did yesterday is the best evidence of what a new material will do on it tomorrow. The first version of that loop can be built today. So the solution is not to bring materials problems to thousands of geniuses in a data center, but to bring the intelligence as well as the observation power of a data center to thousands of labs and plants, close to the equipment and the shop floor.

References

[1] D. Amodei. *Machines of Loving Grace: How AI Could Transform the World for the Better*. 2024. URL: https://www.darioamodei.com/essay/machines-of-loving-grace (accessed Oct. 27, 2026).

[2] A. Merchant et al. "Scaling deep learning for materials discovery". *Nature* 624.7990 (2023), pp. 80–85. DOI: 10.1038/s41586-023-06735-9.

[3] N. J. Szymanski et al. "An autonomous laboratory for the accelerated synthesis of inorganic materials". *Nature* 624.7990 (2023), pp. 86–91. DOI: 10.1038/s41586-023-06734-w.

[4] C. Zeni et al. "A generative model for inorganic materials design". *Nature* 639.8055 (2025), pp. 624–632. DOI: 10.1038/s41586-025-08628-5.

[5] Engineering ToolBox. *Metals and Alloys: Bulk Modulus of Elasticity*. Steel, various grades: 156–165 GPa. 2008. URL: https://www.engineeringtoolbox.com/bulk-modulus-metals-d_1351.html (accessed Nov. 6, 2026).

[6] A. K. Cheetham and R. Seshadri. "Artificial Intelligence Driving Materials Discovery? Perspective on the Article: Scaling Deep Learning for Materials Discovery". *Chemistry of Materials* 36.8 (2024), pp. 3490–3495. DOI: 10.1021/acs.chemmater.4c00643.

[7] J. Leeman et al. "Challenges in High-Throughput Inorganic Materials Prediction and Autonomous Synthesis". *PRX Energy* 3.1 (2024), p. 011002. DOI: 10.1103/PRXEnergy.3.011002.

[8] HPCwire. *Lila Sciences nabs $235 million to build scientific superintelligence*. Secondary coverage. 2025. URL: https://www.hpcwire.com/2025/09/15/lila-sciences-nabs-235-million-to-build-scientific-superintelligence/ (accessed Oct. 27, 2026).

[6] Appendix A gives the calculation. McKinsey's sample is 35 successful launches with data from 2006–2010, so it says nothing about programs that died.

[9] J. Bort. *Former OpenAI and DeepMind researchers raise whopping $300M seed to automate science*. TechCrunch. 2025. URL: https://techcrunch.com/2025/09/30/former-openai-and-deepmind-researchers-raise-whopping-300m-seed-to-automate-science (accessed Oct. 27, 2026).

[10] Fierce Biotech. *Lila Sciences adds Nvidia-backed $115M to series A, bringing total haul to $350M*. Secondary coverage. 2025. URL: https://www.fiercebiotech.com/biotech/lila-sciences-adds-115m-series-bringing-total-haul-350m-and-securing-backing-nvidia (accessed Oct. 27, 2026).

[11] K. J. Jenewein et al. *AI-guided high-throughput discovery of iridium- and ruthenium-free palladium-oxide catalysts for durable acidic oxygen evolution*. Preprint, Lila Sciences. 2026. arXiv: 2609.30133 [cond-mat.mtrl-sci].

[12] L. Zhou et al. "High Throughput Discovery of Solar Fuels Photoanodes in the CuO–V_2O_5 System". *Advanced Energy Materials* 5.22 (2015), p. 1500968. DOI: 10.1002/aenm.201500968.

[13] M. Abolhasani and E. Kumacheva. "The rise of self-driving labs in chemical and materials sciences". *Nature Synthesis* 2.6 (2023), pp. 483–492. DOI: 10.1038/s44160-022-00231-0.

[14] G. Tom et al. "Self-Driving Laboratories for Chemistry and Materials Science". *Chemical Reviews* 124.16 (2024), pp. 9633–9732. DOI: 10.1021/acs.chemrev.4c00055.

[15] U.S. Department of Energy, Hydrogen and Fuel Cell Technologies Office. *Technical Targets for Proton Exchange Membrane Electrolysis*. Stack targets: 3.0 A/cm^2; lifetime 40,000 h (2022 status), 80,000 h (2026 and ultimate). URL: https://www.energy.gov/eere/fuelcells/technical-targets-proton-exchange-membrane-electrolysis (accessed Oct. 27, 2026).

[16] J. Fan et al. "Bridging the gap between highly active oxygen reduction reaction catalysts and effective catalyst layers for proton exchange membrane fuel cells". *Nature Energy* 6.5 (2021), pp. 475–486. DOI: 10.1038/s41560-021-00824-7.

[17] J. Knöppel et al. "On the limitations in assessing stability of oxygen evolution catalysts using aqueous model electrochemical cells". *Nature Communications* 12 (2021), p. 2231. DOI: 10.1038/s41467-021-22296-9.

[18] R. O. Ritchie. "The conflicts between strength and toughness". *Nature Materials* 10.11 (2011), pp. 817–822. DOI: 10.1038/nmat3115.

[19] D. Sun et al. "Why 90% of clinical drug development fails and how to improve it?" *Acta Pharmaceutica Sinica B* 12.7 (2022). The 90% refers to candidates that have entered Phase I trials, pp. 3049–3062. DOI: 10.1016/j.apsb.2022.02.002.

[20] M. K. P. Jayatunga et al. "How successful are AI-discovered drugs in clinical trials? A first analysis and emerging lessons". *Drug Discovery Today* 29.6 (2024). Phase I success of 80–90% for AI-discovered molecules against 40–65% historically; Phase II about 40%, on a limited sample, comparable to historical averages, p. 104009. DOI: 10.1016/j.drudis.2024.104009.

[21] National Science and Technology Council. *Materials Genome Initiative for Global Competitiveness*. White paper. Washington, DC: Executive Office of the President, 2011.

[22] National Research Council. *Accelerating Technology Transition: Bridging the Valley of Death for Materials and Processes in Defense Systems*. Report. Washington, DC: The National Academies Press, 2004. DOI: 10.17226/11108.

[23] O. T. Unke et al. "Machine Learning Force Fields". *Chemical Reviews* 121.16 (2021), pp. 10142–10186. DOI: 10.1021/acs.chemrev.0c01111.

[24] J. Fish, G. J. Wagner, and S. Keten. "Mesoscopic and multiscale modelling in materials". *Nature Materials* 20.6 (2021), pp. 774–786. DOI: 10.1038/s41563-020-00913-0.

[25] P. Friederich et al. “Machine-learned potentials for next-generation matter simulations”. *Nature Materials* 20.6 (2021), pp. 750–761. DOI: 10.1038/s41563-020-0777-6.

[26] J. Riebesell et al. “A framework to evaluate machine learning crystal stability predictions”. *Nature Machine Intelligence* 7.6 (2025), pp. 836–847. DOI: 10.1038/s42256-025-01055-1.

[27] W. Jia et al. “Pushing the Limit of Molecular Dynamics with Ab Initio Accuracy to 100 Million Atoms with Machine Learning”. In: *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. 2020, pp. 1–14. DOI: 10.1109/SC41405.2020.00009.

[28] L. Barroso-Luque et al. “The Open Materials 2024 (OMat24) inorganic materials dataset and models”. *Nature Computational Science* 6.6 (2026). Preprint arXiv:2410.12771: over 110 million DFT calculations, pp. 642–652. DOI: 10.1038/s43588-026-00996-w.

[29] B. Meredig et al. “Can machine learning identify the next high-temperature superconductor? Examining extrapolation performance for materials discovery”. *Molecular Systems Design & Engineering* 3.5 (2018), pp. 819–825. DOI: 10.1039/C8ME00012C.

[30] K. A. Severson et al. “Data-driven prediction of battery cycle life before capacity degradation”. *Nature Energy* 4.5 (2019), pp. 383–391. DOI: 10.1038/s41560-019-0356-8.

[31] P. M. Attia et al. “Closed-loop optimization of fast-charging protocols for batteries with machine learning”. *Nature* 578.7795 (2020), pp. 397–402. DOI: 10.1038/s41586-020-1994-5.

[32] J. Grabowski. *QuesTek Innovations LLC: presentation to the Steel Research Group meeting, Northwestern University*. Slides. States 12 employees; lists the AMS and MMPDS designations of the four Ferrium steels; reports the A-basis tensile minimum for S53 predicted from three heats by the Accelerated Insertion of Materials method. 2014. URL: https://chimad.northwestern.edu/docs/SRG2014/SRG2014_Grabowski.pdf (accessed Oct. 28, 2026).

[33] “Ferrium C64: A 21st Century Steel Takes Flight”. *Advanced Materials & Processes* 179.7 (2021), p. 19. URL: https://static.asminternational.org/amp/202110/23/ (accessed Oct. 27, 2026).

[34] G. B. Olson and C. J. Kuehmann. “Materials genomics: From CALPHAD to flight”. *Scripta Materialia* 70 (2014), pp. 25–30. DOI: 10.1016/j.scriptamat.2013.08.032.

[35] S. Noy and W. Zhang. “Experimental evidence on the productivity effects of generative artificial intelligence”. *Science* 381.6654 (2023), pp. 187–192. DOI: 10.1126/science.adh2586.

[36] E. Brynjolfsson, D. Li, and L. Raymond. “Generative AI at Work”. *The Quarterly Journal of Economics* 140.2 (2025), pp. 889–942. DOI: 10.1093/qje/qjae044.

[37] A. Humlum and E. Vestergaard. *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. NBER Working Paper 33777. Revised March 2026; formerly titled “Large Language Models, Small Labor Market Effects”. National Bureau of Economic Research, 2025. DOI: 10.3386/w33777.

[38] A. Bick, A. Blandin, and D. J. Deming. *The Rapid Adoption of Generative AI*. NBER Working Paper 32966. Time savings of 5.4% of hours among users and 1.4% across all workers. National Bureau of Economic Research, 2024. DOI: 10.3386/w32966.

[39] I. Yotzov, J. M. Barrero, N. Bloom, et al. *Firm Data on AI*. NBER Working Paper 34836. Revised October 2026. National Bureau of Economic Research, 2026. DOI: 10.3386/w34836.

[40] J. W. Scannell and J. Bosley. “When Quality Beats Quantity: Decision Theory, Drug Discovery, and the Reproducibility Crisis”. *PLOS ONE* 11.2 (2016), e0147215. DOI: 10.1371/journal.pone.0147215.

[41] J. W. Scannell et al. “Predictive validity in drug discovery: what it is, why it matters and how to improve it”. *Nature Reviews Drug Discovery* 21.12 (2022), pp. 915–931. DOI: 10.1038/s41573-022-00552-x.

[42] R. S. Bohacek, C. McMartin, and W. C. Guida. “The Art and Practice of Structure-Based Drug Design: A Molecular Modeling Perspective”. *Medicinal Research Reviews* 16.1 (1996), pp. 3–50. DOI: 10.1002/(SICI)1098-1128(199601)16:1<3::AID-MED1>3.0.CO;2-6.

[43] H. M. Cobb. *The History of Stainless Steel*. Materials Park, OH: ASM International, 2010.

[44] R. J. Plunkett. “The History of Polytetrafluoroethylene: Discovery and Development”. In: *High Performance Polymers: Their Origin and Development*. Ed. by R. B. Seymour and G. S. Kirshenbaum. New York: Elsevier, 1986, pp. 261–266. DOI: 10.1007/978-94-011-7073-4_25.

[45] D. Tanner, J. A. Fitzgerald, and B. R. Phillips. “The Kevlar Story—an Advanced Materials Case Study”. *Angewandte Chemie International Edition in English* 28.5 (1989), pp. 649–654. DOI: 10.1002/anie.198906491.

[46] M. Appl. “The Haber–Bosch Heritage: The Ammonia Production Technology”. In: *50th Anniversary of the IFA Technical Conference* (Sevilla, Spain, Sept. 25–26, 1997). International Fertilizer Industry Association, 1997. URL: https://www.fertilizer.org/wp-content/uploads/2023/01/HABER.pdf (accessed Oct. 27, 2026).

[47] C. Mead and L. Conway. *Introduction to VLSI Systems*. Reading, MA: Addison-Wesley, 1980.

[48] Taiwan Semiconductor Manufacturing Company. *Company Profile*. From the 1997 annual report: founded in 1987, the first pure-play integrated-circuit foundry. 1998. URL: https://investor.tsmc.com/static/annualReports/1997/company1.html (accessed Oct. 28, 2026).

[49] AAII Staff. *Which Is a Better Investment, Cadence Design Systems Inc or Synopsys, Inc. Stock?* Market capitalizations as of 3 November 2026: Cadence $88.8 billion, Synopsys $88.7 billion. American Association of Individual Investors. 2026. URL: https://www.aaii.com/investingideas/article/26731-which-is-a-better-investment-cadence-design-systems-inc-or-synopsys-inc-stock (accessed Nov. 4, 2026).

[50] L. Mwogah. *Largest Companies by Market Cap in 2026: Nvidia Leads the Global Rankings*. Secondary coverage. Market data as of 5 October 2026: NVIDIA about $5.65 trillion, TSMC about $2.45 trillion. InvestingCube. 2026. URL: https://www.investingcube.com/deep-analysis/largest-companies-by-market-cap-in-2026-nvidia-leads-the-global-rankings/ (accessed Oct. 28, 2026).

[51] M. Miremadi, C. Musso, and J. Oxgaard. *Chemical innovation: An investment for the ages*. 35 successful launches with data from 2006–2010; figures taken from the project’s research notes. McKinsey on Chemicals. 2013. URL: https://www.mckinsey.com/~/media/mckinsey/dotcom/client_service/chemicals/pdfs/chemical_innovation_an_investment_for_the_ages.ashx.

[52] W. Isaacson. *Steve Jobs*. Recounts how Jobs persuaded Corning’s chief executive, Wendell Weeks, to make Gorilla Glass for the first iPhone within six months. New York: Simon & Schuster, 2011.

[53] C. Westerheide. *Sila recauda $300 millones para ampliar la producción de ánodos de silicio en EE. UU.* Secondary coverage, in Spanish. Moses Lake plant in operation since autumn 2025. electrive. 2026. URL: https://www.electrive.com/es/2026/07/22/sila-raises-300-million-to-expand-us-silicon-anode-production/ (accessed Oct. 28, 2026).

[54] Batteries News. *Group14 Accelerates EV-Scale Production of Breakthrough Silicon Battery Material in South Korea*. Secondary coverage of a 12 March 2026 press release; Moses Lake plant nearing completion. 2026. URL: https://batteriesnews.com/group14-accelerates-ev-scale-production-of-breakthrough-silicon-battery-material-in-south-korea/ (accessed Oct. 28, 2026).

[55] Society of Chemical Manufacturers & Affiliates. *Specialty Chemistry*. Defines a toll manufacturer as a company that sells its manufacturing capabilities and charges the customer a toll fee for the use of its manufacturing assets. URL: https://www.socma.org/about/specialty-chemistry/ (accessed Nov. 4, 2026).

[56] IEEE Spectrum. *Chip Hall of Fame: Xilinx XC2064 FPGA*. The XC2064, Xilinx's first product, came out in 1985. 2017. URL: https://spectrum.ieee.org/chip-hall-of-fame-xilinx-xc2064-fpga (accessed Nov. 4, 2026).

[57] K. Shirriff. *Reverse-engineering the first FPGA chip, the XC2064*. Blog post. "Xilinx was one of the first fabless semiconductor companies"; it used Seiko Epson's semiconductor division for its initial fab. 2020. URL: https://www.righto.com/2020/09/reverse-engineering-first-fpga-chip.html (accessed Nov. 4, 2026).

[58] Lawrence Berkeley National Laboratory, Computing Sciences. *Machine Learning Software that Enhances Molecular Dynamics Modeling Wins Gordon Bell Prize*. Lin Lin on the reach of DFT-based molecular dynamics: hundreds of atoms, up to a few picoseconds per day. 2020. URL: https://cs-newsarchive.lbl.gov/news/2020/machine-learning-software-that-enhances-molecular-dynamics-modeling-wins-gordon-bell-prize/ (accessed Nov. 2, 2026).

[59] J. G. Bednorz and K. A. Müller. "Possible high T_c superconductivity in the Ba–La–Cu–O system". *Zeitschrift für Physik B Condensed Matter* 64.2 (1986), pp. 189–193. DOI: 10.1007/BF01303701.

[60] N. Singh. "Leading theories of the cuprate superconductivity: A critique". *Physica C: Superconductivity and its Applications* 580 (2021), p. 1353782. DOI: 10.1016/j.physc.2020.1353782.

[61] N. Singh. "40 years of cuprate high-T_c superconductors: A perspective on theories". *Physica C: Superconductivity and its Applications* 651 (2026), p. 1354957. DOI: 10.1016/j.physc.2026.1354957.

[62] International Council for Harmonisation. *Pharmaceutical Development Q8(R2)*. ICH harmonised guideline, current Step 4 version of August 2009, as published by the European Medicines Agency. "Working within the design space is not considered as a change. Movement out of the design space is considered to be a change and would normally initiate a regulatory post approval change process." 2009. URL: https://www.ema.europa.eu/en/ich-q8-r2-pharmaceutical-development-scientific-guideline (accessed Nov. 4, 2026).

[63] M. Allsup. *How to monetize AI in the battery sector*. Secondary coverage. Latitude Media. 2024. URL: https://latitudemedia.com/news/how-to-monetize-ai-in-the-battery-sector (accessed Oct. 28, 2026).

[64] J. Manyika and M. Chui. *By 2025, Internet of things applications could have $11 trillion impact*. By two McKinsey Global Institute researchers: a modern oil platform has over 30,000 sensors, and less than 1 percent of their data is used for decision-making. Fortune. 2015. URL: https://fortune.com/2015/07/22/mckinsey-internet-of-things/ (accessed Nov. 4, 2026).

[65] NVIDIA. *Our History: Innovations Over the Years*. Founded 5 April 1993. URL: https://www.nvidia.com/en-us/about-nvidia/corporate-timeline/ (accessed Oct. 28, 2026).

[66] J. Kahn. *Exclusive: Orbital Industries, startup using AI to discover exotic new materials, raises $50 million Series B funding round*. The coolant family is in qualification with major chip providers and designed to ship with the next generation of GPUs in 2027; the chief executive says it would be the first AI-designed molecule on the commercial market in any industry. Fortune. 2026. URL: https://fortune.com/2026/05/28/exclusive-orbital-industries-raises-50-million-series-b-funding-round-ai-to-discover-exotic-new-materials/ (accessed Nov. 4, 2026).

Appendix A: The model math

The article uses three small calculations. Their inputs are assumptions or my judgment, and none has been fitted to data, so they should be read as illustrations.

Validity against throughput

This calculation follows Scannell and Bosley [40]. Each candidate has a true value Y, its performance in the field, and a lab score X. Both are standard normal, and their correlation ρ is the predictive validity of the lab test. A candidate works in the field if Y is above the threshold $y^* = \Phi^{-1}(1-\pi)$, where Φ is the standard normal distribution function and π is the fraction of candidates that work. The lab tests N candidates and advances the one with the highest score, $X_{\max}$. The chance that this candidate works is

$$P(\rho, N, \pi) = \mathrm{E}\left[1 - \Phi\left(\frac{y^* - \rho X_{\max}}{\sqrt{1-\rho^2}}\right)\right].$$

The expectation is over $X_{\max}$, the largest of N standard normal draws, and we estimate it from 400,000 samples of $X_{\max}$. The space of possible candidates is far larger than any lab can test, so we treat the N tested candidates as random draws from it. Table A1 gives this chance where one candidate in a million works ($\pi = 10^{-6}$).

Table A1: Chance in 10,000 that the advanced candidate works, for $\pi = 10^{-6}$.

Validity ρ	$N = 10^6$	$N = 10^7$	$N = 10^8$
0.30	2.9	4.7	7.4
0.33	4.5	7.6	12.2
0.40	11.7	21.4	36.5

At a validity of 0.30, each tenfold increase in tests multiplies the chance by about 1.6. Raising the validity by ten percent, to 0.33, multiplies it by 1.55 at a million tests, about what ten times more tests do (1.64). Raising the validity by a third, to 0.40, multiplies it by 4.0, compared with 2.5 for a hundred times more tests. To match the 0.40 test at a million candidates, a test of validity 0.30 would need about 1.4 billion. The validity of 0.30 and the rate of one in a million are assumptions, and the model has one test and one property.

The clock

Stage i takes a share s_i of today's time, and we assume AI makes it g_i times faster. The time left, as a fraction of today's, is then

$$\frac{T_{\text{new}}}{T_{\text{today}}} = \sum_i \frac{s_i}{g_i}.$$

Table A2 applies this to the midpoints of the shares and gains in the stage table. Physical understanding has no share of its own, so it is not in the sum.

Table A2: The clock at the stage table's midpoints, in percent of today's time.

Stage	Share	Gain	Time left	Time saved
Candidate generation	3.0	10×	0.3	2.7
Lab synthesis and screening	7.5	10×	0.8	6.8
Design validation	15.0	2.5×	6.0	9.0
Scale-up	25.0	2×	12.5	12.5
Qualification	30.0	1.5×	20.0	10.0
Adoption	20.0	1.35×	14.8	5.2
Total	100.5		54.4	46.1

The time left is 54.4 percent of today's, a factor of 1.84, so a program of fifteen years would take about eight. The low and high ends of the gains give factors of 1.73 and 1.93. Halving scale-up saves 12.5 points of the time, and a lab ten times faster saves 6.8. The midpoints of the shares sum to 100.5 percent, and the factor is 100 divided by the time left. The stage table describes the gain in candidate generation as several-fold, and the calculation uses 10. With a gain of 3 there, the factor would be 1.82. A set of shares for a short ladder (3, 12, 25, 30, 15 and 15 percent, in the table's order) gives 2.10.

The return on a program

A program spends a total S at a level rate over D years of development, and then earns a level return a a year for L years. Its rate of return r is the rate at which the discounted return equals the discounted spending:

$$a\, e^{-\delta D}\, \frac{1 - e^{-\delta L}}{\delta} = \frac{S}{D}\, \frac{1 - e^{-\delta D}}{\delta}, \qquad \delta = \ln(1 + r).$$

We set D to 14 years, McKinsey's average for new technology in a new market [51], and L to 15 years, and we choose a so that r is 10 percent. Dividing D by the clock's factor of 1.84 gives 7.6 years. With the same total spending, r would rise to 14.1 percent. If spending falls in proportion to the time, r would rise to 21.8 percent. With a product life of 10 years the two results would be 14.8 and 23.6 percent, and with 25 years they would be 13.4 and 20.2 percent.