Table of Contents

The Latency Collapse

Parts I and II — Research Archive Edition

> **Archive note (September 9, 2026):** Historical dates are research estimates rather than point-precise facts. The future timeline is scenario scaffolding, not a calibrated forecast. The September 8, 2026 Navier-Stokes result is described as a proposed solution pending independent scrutiny.

The Latency Collapse

Part I — How General-Purpose Technologies Become Culture

A child sitting in a chat room in 1996 did not think she was participating in the "information revolution." She was talking to people.

Her parents, meanwhile, might still have been asking what a modem was, why anyone would want to talk to strangers through a computer, what the Internet was actually *for*, and whether this strange new thing was mostly hype. The child did not need an answer. She was already inside it.

A child talking to an AI in 2026 occupies almost exactly the same cultural position.

Adults are asking what GPT means. They are debating whether artificial general intelligence is possible, whether AI is merely autocomplete, whether people should use it to write, whether it will destroy jobs, whether it can truly "reason." Children increasingly encounter the same technology less as an object requiring philosophical justification than as part of the environment. You ask it things. You make things with it. It is there.

This is one of the peculiar characteristics of a general-purpose technology: eventually it stops feeling like technology.

It becomes part of what the world is made of.

When technology disappears

Economists use *general-purpose technology* for a special class of innovations. These are not merely useful inventions. They are technologies that continue improving, spread across large portions of the economy, acquire uses their inventors could not have anticipated, and provoke complementary innovations around themselves. Lipsey, Carlaw and Bekar identified 24 historical "transforming" GPTs, beginning with the domestication of plants and animals and extending through the internet, biotechnology and nanotechnology. [P1-01]

The list is strange at first glance. It contains objects such as the wheel and computer, processes such as writing and biotechnology, and organizational systems such as the factory and lean production.

That is a feature rather than a defect.

The commonality is not physical form. It is **generativity**. A GPT opens a design space in which thousands of other things become possible.

Domesticated plants did not merely produce more food. They made dense settlements, stored surpluses, occupational specialization, new property arrangements and eventually states possible. Domesticated animals became food, clothing, transport, traction and power. Writing became accounting, law, scripture, literature, taxation, bureaucracy and memory outside the brain. Electricity became light, motion, heat, communication, computation and eventually the energy substrate of nearly every other modern technology.

A successful GPT therefore undergoes a conceptual disappearance. We stop discussing the technology itself and begin discussing the world built on top of it.

Nobody advertises a company as an *electrified business*. We rarely distinguish a *computerized corporation*. Increasingly, even “online business” sounds redundant. The qualifier disappears when the technology has become the assumed substrate.

AI will probably become culturally pervasive in the fullest sense at approximately the moment when the word *AI* begins to disappear in the same way.

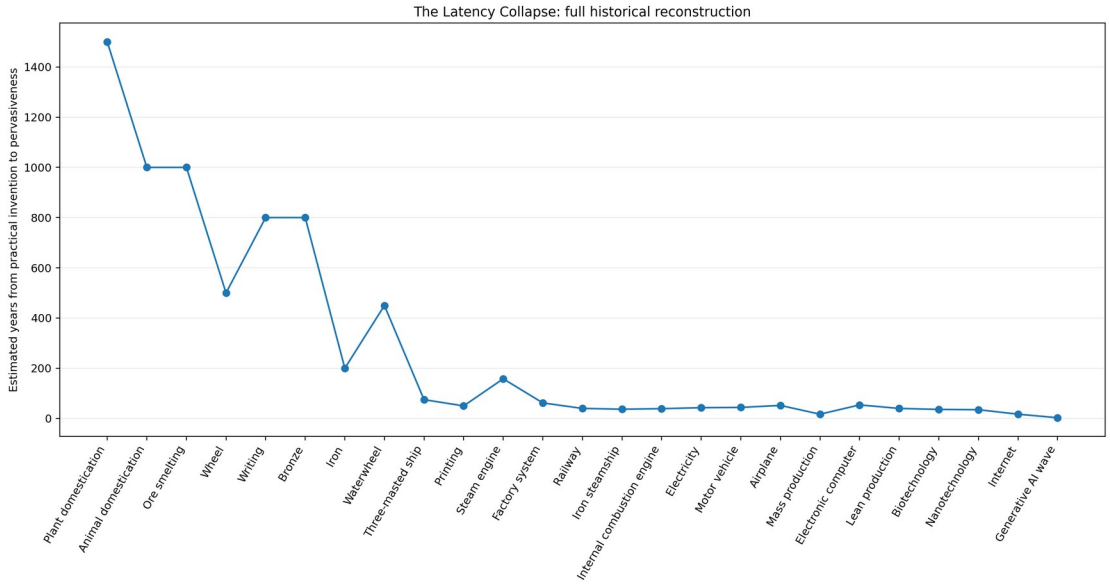

The full historical latency series

The first diffusion regime: technology moved at human speed

For most of human history, invention and cultural pervasiveness could be separated by centuries or millennia.

Plant domestication began in Western Asia during the ninth millennium BCE; animals followed contemporaneously or somewhat later. Their diffusion across adjoining regions unfolded over subsequent millennia. [P1-02] Metallurgy shows the same pattern. Current archaeological evidence puts early copper smelting around the fifth millennium BCE in the Balkans, while mature metallurgical systems spread much more gradually. [P1-03]

The wheel appears in wheeled vehicles in the fourth millennium BCE and was common enough by around 3000 BCE to represent a transformed transportation system. Writing began as specialized administrative recording around the late fourth millennium BCE, but by the middle of the third millennium cuneiform was being used for economic, political, religious, literary and scholarly purposes — roughly the transition from specialist technology to cultural substrate. [P1-04]

The important constraint was not human intelligence. Ancient people were perfectly capable of recognizing useful technologies.

The constraint was the **network**.

Knowledge moved with bodies. Techniques traveled through migration, apprenticeship, conquest, marriage and trade. Materials moved at animal or sailing speed. An inventor could not upload instructions. There were no technical journals, standardized education systems or globally interoperable supply chains. A discovery in one valley could remain effectively invisible to people a few hundred kilometers away.

Even the history of the waterwheel illustrates this. Older histories often claimed that Romans knew the technology but scarcely used it until medieval Europe. Archaeology has overturned much of that view: water power was substantially more common during the Roman period than the literary record once suggested. Yet even this remarkably valuable way of converting natural energy into mechanical work diffused on a scale measured in centuries. [P1-05]

The bottleneck was transmission.

Printing changes the equation

Then something qualitatively different happens.

A securely dated depiction of a three-masted ship exists from the early fifteenth century; by the late fifteenth century the three-master had become a common European vessel type. Printing is more important still.

Movable-type printing in Europe moves from the Gutenberg era around 1450 to a continental network of presses, books and printers around 1500 in roughly half a century.

But the great importance of printing is not simply that *printing spread quickly*.

It is that printing changed the speed at which **other ideas could spread**.

This is the beginning of a recursive process.

A better plow changes farming. A printing press changes farming, astronomy, navigation, mathematics, medicine, theology, engineering — and the ability of people working in all of those fields to communicate improvements to one another.

Technology had begun improving the distribution system for technology.

That distinction becomes increasingly important from this point forward.

The industrial age settles into one human generation

The Industrial Revolution does not produce a simple exponential collapse in adoption time.

Instead, something stranger appears in the data: a large group of transformative technologies converge on a diffusion interval of roughly **one generation**.

The first commercially successful Newcomen steam engine dates to 1712. If that is treated as the starting point and the steam-dominated industrial economy of the later nineteenth century as the endpoint, the transition takes roughly a century and a half. Starting instead from Watt's much more efficient commercial engines in the 1770s cuts the interval dramatically. Steam is therefore one of the clearest demonstrations that "date of invention" is not an objective point but a modeling decision.

The factory system provides a cleaner example. Arkwright's Cromford mill of 1771 became a blueprint copied elsewhere. [P1-06] The railway is cleaner still: passenger rail emerged in northern England in the 1820s and 1830s, and by 1870 Britain was recording hundreds of millions of passenger journeys annually. [P1-07]

Practical internal combustion arrives with Otto's four-stroke engine in 1876. Commercial electrical systems begin transforming cities in the 1880s. Benz's automobile follows in 1886, and by 1930 more than half of American families owned a car. [P1-08] Powered flight begins in 1903; by the mid-1950s the era of mass air travel was well under way. [P1-09]

These are astonishingly different technologies, yet their social assimilation occurs on remarkably similar timescales.

That is because the constraint has shifted.

The problem is no longer primarily getting the *idea* from one place to another. The problem is rebuilding the physical and institutional environment around it.

Electricity requires generation stations, transmission systems, wiring, motors, standards, financing, utilities and regulation. Cars require factories, fuel refining, filling stations, paved roads, traffic law, insurance, repair shops, parking and an entire redesign of cities. Aviation requires aircraft factories, airports, navigation, air-traffic systems, regulation and trained personnel.

A GPT is not simply adopted.

Society has to rearrange itself around the GPT.

And for much of industrial history, rearranging society seems to have required approximately one human generation.

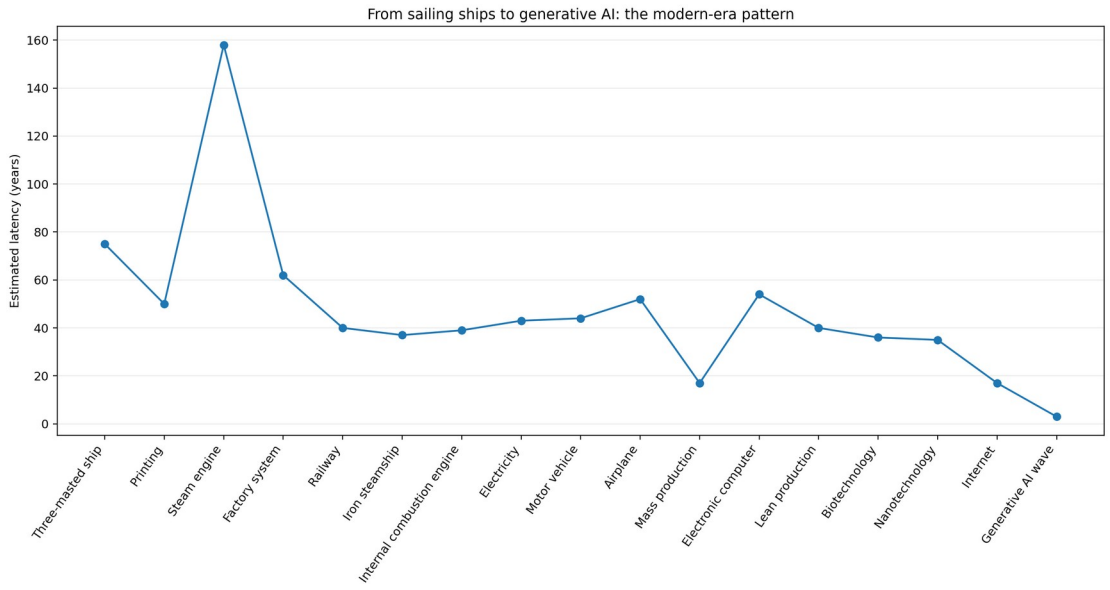

Modern-era latency zoom

Mass production begins to squeeze the interval

Ford's moving automobile assembly line began operating in 1913. Production increased while the price of the Model T fell, and Ford ultimately produced millions of cars using the system. [P1-10]

This mattered beyond automobiles. The moving assembly line was an organizational technology: decomposition of complex production into standardized operations, interchangeable components, synchronized flow and enormous scale.

That makes mass production another recursively important GPT. It reduced the cost and time required to manufacture the physical artifacts through which subsequent technologies would diffuse.

The technological system was beginning to stack.

Printing accelerated information.

Steam and rail accelerated matter and people.

Electricity created a standardized energy substrate.

Mass production accelerated replication.

Each new layer inherited the layers beneath it.

Computers reveal why the starting date matters

The computer appears at first to violate the acceleration story.

If we begin with ENIAC in 1946 and use majority household computer penetration as the endpoint, the United States does not cross that threshold until 2000, when Census data show 51 percent of households owning at least one computer. That produces a 54-year diffusion interval. [P1-11]

But this comparison hides the technology's internal evolution.

A room-sized electronic computer in 1946 and a personal computer in 1981 are members of the same technological lineage, but not equivalent consumer propositions. If the IBM-PC era is treated as the beginning of the mass-market implementation, majority penetration arrives in roughly 19 years rather than 54.

The same problem recurs with almost every mature GPT.

Was steam invented with Newcomen or Watt?

Was the internet invented with ARPANET in 1969, the TCP/IP internet in 1983, the World Wide Web around 1989–91, or the commercial web of the mid-1990s?

Was contemporary AI invented at Dartmouth in 1956, with modern deep learning, with the transformer in 2017, or with the mass-market generative-AI interface in 2022?

There is no neutral answer.

But instead of weakening the argument, this reveals something more important.

Technological gestation and technological deployment are different clocks.

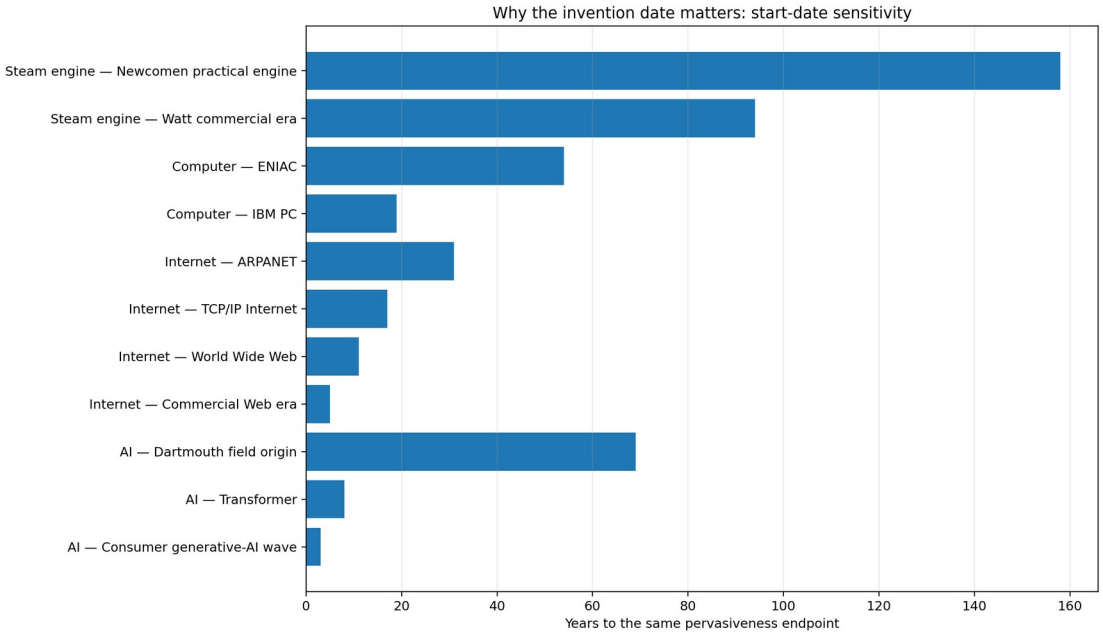

Start-date sensitivity

The Internet attacks the deployment clock

On January 1, 1983, ARPANET completed its transition to TCP/IP — the technical architecture from which the modern Internet developed. [P1-12] By 2000, Pew found that a majority of American adults were using the Internet. [P1-13]

That gives roughly seventeen years from the TCP/IP platform to majority adult adoption.

Begin with ARPANET and the interval is approximately 31 years.

Begin with the World Wide Web and it is about eleven.

Begin with the commercial consumer Internet of the mid-1990s and it is closer to five.

Those are not contradictory measurements. They measure different layers of the same technological transformation.

And the shortest interval is historically the most consequential, because something unprecedented had happened: **the new GPT had created a near-zero-marginal-cost global distribution system for software and information.**

A railway company entering another city had to lay track.

An electric utility entering another town had to string wire.

A car industry entering another country needed factories, roads and fuel distribution.

An Internet service entering another country often needed only servers, software and an already-connected user.

The physical installation bottleneck was beginning to disappear.

AI is arriving through infrastructure built by its predecessor

This is where the analogy between the 1990s chat-room child and today's AI-native child becomes more than a metaphor.

Artificial intelligence itself is old. The academic field is conventionally traced to the Dartmouth project of 1956. If we measure from that intellectual origin to widespread AI use today, the latency is about seventy years — not fast at all.

But today's **deployment wave** is different.

Stanford's 2026 AI Index reports that generative AI reached 53 percent population adoption within three years, faster than the personal computer or Internet in its comparison. [P1-14]

And if the threshold is merely geographic availability rather than deep cultural adoption, the number becomes almost absurd. WIPO reports that generative AI, exemplified by ChatGPT's November 2022 release, had users in virtually every country within days, attributing this speed to a ready-made global digital infrastructure. [P1-17]

This is the crucial point.

AI did not have to construct its own distribution network.

The Internet had already built it.

Computers were already ubiquitous.

Cloud data centers already existed.

Billions of people already carried networked computers in their pockets.

Online payment systems already worked.

Software distribution was already instantaneous.

Digital identity and authentication already existed.

An AI company could therefore release a new cognitive capability and, in principle, make the same underlying system available to billions of people essentially at once.

The invention time may still be long.

The **deployment latency has collapsed**.

An independent dataset sees the same thing

This pattern is not merely an artifact of cherry-picking famous inventions.

A 2026 WIPO study reconstructed diffusion for **31 technologies across 139 countries over two centuries**. Using a standardized measure — the lag between technological invention and recorded adoption in a country — the researchers found average adoption lags falling from

approximately **50 years for technologies before 1950 to 15 years for technologies after 2000**. [P1-15, P1-16]

That measure is deliberately less demanding than the “cultural pervasiveness” threshold used here. A technology can appear in a country long before it becomes normal there.

But that is precisely what makes the result useful.

Two different measures reveal two different layers of acceleration:

Geographic availability is approaching zero. Cultural and institutional transformation is still slower.

This distinction explains much of the confusion surrounding AI today.

AI can be *everywhere* before society has reorganized around it.

The recursive stack

The long historical series therefore suggests something more interesting than “technology is speeding up.”

General-purpose technologies increasingly form a **recursive diffusion stack**.

Plant and animal domestication increased the surplus available for specialization.

Writing externalized memory.

Printing replicated memory cheaply.

Steam increased mechanical power.

Rail and steamships compressed physical distance.

Electricity supplied fungible power almost anywhere.

Mass production replicated complex physical artifacts.

Computers made processes programmable.

Lean production made industrial adaptation more responsive.

Biotechnology made biological systems engineerable.

Nanotechnology pushed manipulation toward molecular scales.

The Internet made digital information globally distributable.

AI makes portions of cognition — translation, explanation, programming, synthesis, search, design and increasingly action — available as an on-demand computational resource.

And that last step is different in kind.

Printing did not read the books it reproduced.

Railways did not design better locomotives.

Electricity did not teach engineers electrical engineering.

The Internet moved information but did not, by itself, understand the information being moved.

AI can participate directly in the process through which people understand, adapt and implement new technologies.

That creates a possible new feedback loop:

technology -> cheaper cognition -> faster technological adoption -> better technology -> still cheaper cognition.

WIPO explicitly discusses digital technologies and modern AI systems as mechanisms that may reduce the costs of finding, interpreting and adapting technical knowledge. [P1-15]

If that effect proves durable, AI may not merely diffuse quickly. It may reduce the diffusion time of whatever comes after it.

That would be historically unusual.

What happens when the cost approaches zero?

When a GPT becomes dramatically cheaper, people do not simply consume more of the exact thing they consumed before. They reorganize production around the new cost structure.

Cheap mechanical power did not merely let factories perform existing hand labor faster. It changed where factories were built and how work was organized.

Cheap electricity did not merely replace candles. It eventually reorganized factories, homes, cities, communication and entertainment.

Cheap computation did not merely automate arithmetic. It made software, digital finance, search engines, videogames, global logistics and social networks possible.

So if useful cognitive work becomes dramatically cheaper, the long-run result should not be imagined simply as “AI does the job a person used to do.”

Some tasks will indeed substitute toward machines. But other human abilities become complements, new products become economically possible, and scarcity migrates elsewhere.

When text is cheap, judgment becomes scarcer.

When code is cheap, knowing what should be built matters more.

When content is abundant, trusted attention becomes scarcer.

When synthetic interaction is abundant, authentic human relationship may become more valuable.

When analysis is cheap, accountability for decisions remains expensive.

When digital intelligence is abundant, energy, land, physical infrastructure, materials and embodied action can become binding constraints.

Capital flows toward the newly productive technology, but value does not simply disappear. It relocates.

This is what GPTs do: they rewrite the map of scarcity.

The children usually arrive first

That brings us back to the child in the chat room.

The adults of the 1990s were often trying to understand the Internet as a discrete object. *What is it for? Why do I need it? Isn't it mostly a toy?*

Those questions made sense only during the transition.

Eventually the Internet ceased to have a single "use." Asking what the Internet was for became comparable to asking what electricity was for.

It was for whatever people subsequently built with it.

AI is approaching the same semantic transition.

A generation growing up with conversational machines will not experience them first as an economic theory, a benchmark score or a philosophical controversy over machine intelligence. Children will discover affordances first. They will develop social norms, habits, abuses, defenses, forms of play and ways of thinking that cannot be predicted from a product demo.

That is almost exactly what happened online.

By the time the adults agree on what the new medium means, the children may already have constructed a culture inside it.

The curve is not exponential. It may be stranger than that.

The deepest lesson from the historical record is therefore not that every GPT spreads faster than the one before it.

The data do not support such a neat story.

Premodern GPTs often required centuries or millennia.

From roughly the early modern period through much of the twentieth century, many transformative technologies converged on something like a one-generation transition: forty years here, fifty there, sometimes much longer when enormous physical infrastructure had to be built.

Then digital distribution changes the governing constraint.

The Internet itself takes years to reach mass adoption, but once present, it makes subsequent digital technologies globally available almost immediately.

Generative AI is the clearest case yet: decades of underlying scientific gestation followed by a mass-market deployment measured in a few years and worldwide initial reach measured in days. [P1-14, P1-17]

So perhaps the quantity that is collapsing is not the time required to invent.

It is the distance between **capability and exposure**.

For most of history, a working invention could exist while almost everyone on Earth remained unaware of it.

Today, a working digital invention can become accessible to much of humanity before the culture has even invented the vocabulary with which to discuss it.

And AI adds one final recursive possibility.

The printing press accelerated the transmission of knowledge.

The Internet accelerated the transmission of essentially all information.

AI may accelerate the **interpretation and application of the information being transmitted**.

If so, we are not simply watching another point on the old technological curve.

We may be watching the mechanism that determines the slope of the next one.

That is why today's children using AI resemble the children in 1990s chat rooms.

Not because AI and the Internet are the same technology.

Because in both cases, the children are already living in a world that the adults are still trying to name.

The Latency Collapse

Part II — When Intelligence Enters Matter

Part I ended with a possibility:

The printing press accelerated the transmission of knowledge.

The Internet accelerated the transmission of essentially all information.

Artificial intelligence may accelerate the interpretation and application of that information.

But that conclusion was incomplete.

Because interpretation is only one remaining bottleneck.

An intelligence can understand how to build a bridge without pouring concrete. It can derive a new alloy without melting metal. It can design a reactor without welding the vessel. It can

propose a drug without synthesizing the molecule. It can discover a better solar cell without manufacturing ten billion of them.

For the first few years of the AI revolution, this distinction allowed us to maintain a comforting conceptual separation.

There was the digital world, where AI could move at extraordinary speed.

And there was “meatspace,” where humans still had to do things.

But what happens when the intelligence gets hands?

What happens when millions, and eventually perhaps billions, of machines can see the physical world, reason about it, manipulate it, repair it and work within it for as long as power and maintenance permit?

What happens when the same intelligence controlling those machines can also do mathematics, write software, design experiments, discover materials and operate laboratories?

At that point, AI stops being merely another general-purpose technology.

It begins to resemble something historically unprecedented:

a general-purpose technology capable of participating in the invention, deployment and reproduction of general-purpose technologies themselves.

That closes the loop begun thousands of years ago.

And once the loop closes, the most important question about technological progress may no longer be how rapidly humans can invent.

It may become:

What bottleneck remains that the system cannot yet automate?

1. The first bottleneck to fall is cognition

Consider what happened during a few months of 2026.

Early AI work on Erdős problems was rightly greeted with skepticism. Some supposed discoveries turned out to be rediscoveries of known mathematics. But the line kept moving. In May 2026, Nature reported that an AI system had disproved an 80-year-old unit-distance conjecture associated with Paul Erdős after a single prompt from OpenAI mathematicians. [P2-01]

Then, on September 8, OpenAI published something more dramatic: a proposed solution to the Navier-Stokes existence and smoothness Millennium Prize Problem, together with a Lean formalization. [P2-02]

This result should be treated carefully. Publication by OpenAI is not the same thing as historical mathematical settlement, and the Clay Mathematics Institute’s acceptance process is

deliberately much stricter. But from the standpoint of technological history, another fact is already important.

Mathematical research has become parallelizable industrial production.

A human mathematician is difficult to copy.

Ten thousand world-class mathematicians cannot be created because a server has spare capacity on Tuesday.

They cannot instantly share memories.

They cannot all work continuously for four days.

They cannot be instantiated, assigned different approaches to a problem, terminated when an approach fails, copied when one succeeds, and merged back together.

Software can.

This creates an entirely different economics of cognition.

The relevant unit of productive intelligence begins moving from the *person* toward something more like the **intelligence-hour**.

And intelligence-hours can be manufactured.

2. From intelligence as talent to intelligence as infrastructure

Human civilization has historically been extraordinarily constrained by the rarity of cognitive talent.

There may be billions of people, but only some fraction become mathematicians.

A smaller fraction specialize in partial differential equations.

A smaller fraction still become exceptional researchers.

Each has one brain.

Each brain has twenty-four hours in a day.

Each researcher must sleep.

Each must learn the relevant literature.

Each develops only a limited number of ideas during a career.

Each eventually dies.

The structure of scientific progress has therefore always been constrained by a brutally simple fact:

thinking is scarce.

That scarcity quietly structures the entire economy.

We pay lawyers because legal reasoning is scarce.

Engineers because engineering judgment is scarce.

Programmers because the ability to transform specifications into functioning software is scarce.

Doctors because diagnostic expertise is scarce.

Scientists because discovering something humanity does not yet know is exceptionally scarce.

Managers because organizational judgment is scarce.

Teachers because individualized explanation and attention are scarce.

The knowledge economy is, in large part, an elaborate system for allocating scarce cognition.

AI begins dissolving that scarcity.

Inside OpenAI, agent runtime crossed an interesting threshold in 2026: by mid-August, the research organization reported using about **3.1 agent-workdays of effort for every human workday**. [P2-03]

Independent measurements point in the same general direction while warning against naive extrapolation. METR's task-completion horizon research tracks how long a human-expert task frontier AI agents can complete with specified reliability, and the historical curve has risen rapidly. [P2-04]

The exact date at which agents can independently execute a week, a month, or a year of expert work matters less than the mechanism.

A profession does not need to disappear all at once.

Instead, its workflow becomes a chain.

Research question.

Literature review.

Data collection.

Analysis.

Code.

Draft.

Critique.

Revision.

Presentation.

A human may initially perform seven of those ten links.

Then five.

Then three.

Eventually the human approves the beginning and checks the end.

Then perhaps the approval itself becomes routine.

Recent economic work formalizes this intuition: adjacent AI-executable steps can connect into longer chains, creating nonlinear productivity gains as capabilities improve. [P2-05]

This is how “AI assists knowledge workers” could become “AI performs knowledge work” much faster than occupational labels suggest.

The lawyer may exist long after most legal production is automated.

The programmer may exist long after most code is machine-generated.

The scientist may remain responsible for a laboratory whose hypotheses, experiment schedules, analyses and papers are overwhelmingly produced by machines.

Job titles can persist after the underlying production function has changed.

3. Then intelligence gets hands

But digital cognition alone encounters a wall.

Atoms.

You can generate one million architectural designs in an hour.

You cannot generate one million buildings in an hour.

You can simulate one million molecules.

Someone must still synthesize the promising ones.

You can design a better electric motor.

Copper must still be drawn.

Steel must still be machined.

Magnets must still be manufactured.

The motor must still be assembled.

This is the next great bottleneck migration.

As cognition becomes cheaper, **physical execution becomes relatively more expensive**.

And so capital flows toward automating physical execution.

This transition has already begun.

The International Federation of Robotics counted approximately **4.66 million industrial robots operating worldwide in 2024**, with 542,076 new installations during that year. [P2-06]

Most are nothing like science-fiction androids: they occupy carefully engineered environments and perform narrow repetitive tasks extraordinarily well.

The significance of humanoid and generalist robotics is that it attacks the opposite problem.

Instead of rebuilding the environment around the machine, build a machine that can operate in environments already built around the human body.

Doors.

Stairs.

Shelves.

Forklifts.

Tools.

Factory stations.

Warehouses.

Kitchens.

Hospitals.

Construction sites.

Our civilization is full of interfaces designed for two arms, two hands, human height and human-scale force.

A sufficiently capable general robot inherits access to all of them.

This is why even today's narrow deployments matter.

At BMW's Spartanburg plant, Figure reports that its Figure 02 robots accumulated **1,250+ hours of runtime, loaded 90,000+ parts, and contributed to production of 30,000+ X3 vehicles**. [P2-07]

Meanwhile robot foundation models such as NVIDIA's GR00T are explicitly pursuing transferable vision-language-action systems: learned control architectures that can accept images, language and robot state and generate actions across multiple tasks and embodiments. [P2-08]

The analogy with language models should be obvious.

Robotics historically required programming each task.

Modern physical AI is trying to produce something closer to a pretrained **general model of action**.

If that works, the economics of robotics changes dramatically.

Installing a new skill may increasingly resemble installing software.

4. When software can operate hardware

This is the moment the Internet analogy breaks.

The Internet achieved extraordinary diffusion because software can be copied nearly for free.

A new AI capability released in California can be available in Virginia, Tokyo, Nairobi and Buenos Aires almost immediately.

Robots cannot be downloaded.

They require motors, actuators, batteries, chips, gears, sensors and metal.

So the physical world retains a manufacturing latency.

But now imagine that the robots participate in manufacturing robots.

Not literally one humanoid alone digging ore, refining silicon and assembling an identical humanoid from scratch.

Industrial civilization is not organized that way.

Instead consider the whole network:

AI designs a better robot.

AI manages procurement.

Autonomous mining equipment extracts materials.

Autonomous trucks move them.

Robot-run factories fabricate components.

Automated warehouses distribute them.

Robots assemble additional robots.

Other robots construct new factories.

AI operates the electrical system supplying them.

AI schedules maintenance across the fleet.

No individual machine is self-replicating.

The industrial ecosystem is.

This would amount to distributed self-replication of productive capital.

And that is qualitatively different from every previous GPT.

A steam engine could power the factory that manufactured additional steam engines, so there is historical precedent for technologies assisting their own reproduction.

But the steam engine did not decide how the factory should be redesigned.

The assembly line did not discover a better alloy.

Electricity did not negotiate a supply contract.

The industrial robot did not traditionally rewrite its own controller after observing a new production problem.

Artificial intelligence potentially connects all of these layers.

The result is an autocatalytic industrial system:

intelligence builds machines -> machines build productive capacity -> productive capacity builds more computation -> computation produces more intelligence.

The loop need not become infinitely fast.

Physics forbids that.

Factories take time to construct.

Mines take time to develop.

Power plants require materials.

Machines wear.

Supply chains encounter bottlenecks.

But the quantity that disappears is different:

the requirement that productive capacity scale in proportion to available human labor.

That would be one of the largest discontinuities in economic history.

5. Science becomes a machine

There is another loop closing at exactly the same time.

For centuries the scientific method has consisted of something approximately like:

observe -> hypothesize -> design experiment -> build experiment -> measure -> analyze -> revise hypothesis -> repeat.

Humans occupied nearly every arrow.

AI first entered at analysis.

Then literature search.

Then prediction.

Then experiment design.

Robotics enters at execution.

At that point the loop can run without a human operating every step.

This is not hypothetical anymore.

Berkeley's A-Lab combined machine learning, literature-derived synthesis recipes, computation and robotic experimentation. During **17 days of continuous autonomous operation it synthesized 36 of 57 target compounds**, while active learning proposed follow-up recipes after failures. [P2-09]

The implication is profound.

The scientific method itself becomes **scalable infrastructure**.

Suppose a human materials scientist can meaningfully investigate ten candidate compositions during some interval.

A robotic laboratory might investigate hundreds.

A network of laboratories might investigate tens of thousands.

Simulations may eliminate most candidates before a physical experiment occurs.

AI learns from every failure.

The next experiment is selected automatically.

The loop continues overnight.

On Christmas.

During a pandemic.

After the scientist retires.

And unlike a human research group, successful reasoning strategies can be copied instantly to every other laboratory in the network.

This transforms one of civilization's slowest processes.

Science has traditionally been **serial**.

Automated science can become massively parallel.

6. The room-temperature-superconductor example

Consider superconductivity.

A practical room-temperature superconductor — meaning not merely room-temperature behavior under extraordinary pressure, but a material that works near ambient conditions, carries useful current, tolerates useful magnetic fields, can be manufactured at scale and is made from economically obtainable materials — would be transformative.

Finding such a thing has traditionally required some mixture of physical insight, luck, computational screening and laboratory experimentation.

Now look at what already exists.

A 2026 AI-accelerated superconductivity project began with more than **1.3 million candidate structures**. Machine-learning models and physics calculations reduced that space to **741 dynamically and thermodynamically stable compounds with DFT-confirmed predicted superconductivity above 5 K**, and researchers experimentally synthesized and confirmed two previously unreported superconductors. [P2-10]

Those were not room-temperature superconductors.

That is not the point.

The point is the **search architecture**.

Imagine successive generations of that system.

A model proposes ten million structures.

Specialized models predict stability.

Other models predict electron behavior.

Quantum calculations examine the survivors.

AI designs synthetic pathways.

Robotic laboratories fabricate them.

Automated characterization measures their properties.

The failures return to the model.

The model updates its search strategy.

Thousands of autonomous facilities conduct variations of the process simultaneously.

Humanity has changed the problem from:

> Can a brilliant materials scientist think of the right compound?

into:

> How much of chemical possibility can our automated civilization economically search?

That is an entirely different question.

And AI is unusually suited to problems with enormous combinatorial search spaces.

7. Fusion illustrates the same transition

Fusion makes the distinction between *scientific solution* and *civilizational deployment* even clearer.

There is no single theorem called “fusion.”

Commercial fusion requires many interlocking achievements:

plasma stability,

magnets,

materials resistant to extreme neutron flux,

fuel handling,

tritium management in some designs,

heat extraction,

maintenance,

manufacturing,

control systems,

economics,

regulation,

power conversion.

Solving one does not solve the others.

But AI can attack many of them simultaneously.

In 2026, Princeton Plasma Physics Laboratory reported that its PACMAN AI framework had been tested on the DIII-D tokamak, with machine-learning controllers making decisions on millisecond timescales and predicting a dangerous instability early enough to intervene. [P2-11]

Google DeepMind has also worked with Commonwealth Fusion Systems on AI approaches for fusion development and plasma behavior. [P2-12]

Again, these facts do **not** mean AI has solved commercial fusion.

They show something more relevant to the century ahead:

AI is entering the feedback loops through which fusion will either be solved or shown to be economically unattractive.

Now add autonomous materials science.

Add automated engineering.

Add robotic construction.

Add automated simulation.

Add thousands of parallel AI research teams.

The distinction between “AI solves fusion” and “humans solve fusion using AI” may eventually become semantically meaningless.

The scientific enterprise itself is becoming hybrid.

Eventually it may become predominantly machine-operated.

8. The abundance stack

Suppose several of these loops succeed.

Not magically.

Not all at once.

But over several decades.

Imagine civilization acquires four increasingly abundant inputs:

cognition

labor

energy

materials engineering

Each amplifies the others.

Cheap cognition improves robot design.

Robots reduce the labor cost of building energy infrastructure.

Cheap energy reduces the operating cost of robots and computation.

Better materials improve batteries, turbines, reactors, chips and robots.

Better computers produce more cognition.

More cognition discovers better materials.

The result is another recursive stack.

Part I described the historical stack:

writing -> printing -> steam -> rail -> electricity -> mass production -> computers -> Internet -> AI.

Part II adds a loop:

AI -> robotics -> autonomous science -> better energy/materials -> more compute -> better AI -> better robotics.

There is no obvious historical equivalent.

The steam engine accelerated industrialization.

Electricity accelerated industrialization.

Computers accelerated information processing.

But none potentially occupied every causal link simultaneously.

AI may.

This is why the transition from digital AI to embodied AI matters so much.

It turns the intelligence revolution into a **production revolution**.

9. The next twenty years: 2026–2046

The first phase is likely to be much less cinematic than science fiction.

There will not necessarily be one morning when everyone discovers that all jobs have vanished.

Instead the economy can develop enormous disparities in productivity.

A small team operating hundreds of agents can begin performing work that previously required a large organization.

Research groups can run far more experiments.

Software production can become increasingly machine-dominated.

Professional services can reorganize around AI-generated work reviewed by progressively fewer humans.

The ILO cautions that AI exposure measures should not be interpreted as direct job-loss forecasts, but its 2026 review also notes that modern capability-based measures place cognitive, analytical, administrative and managerial work among the most exposed categories. [P2-13]

This matters because the new automation wave reaches directly into the occupations industrial civilization spent two centuries educating people toward.

Previously, technological automation tended to threaten physical routine work first.

Now accountants, analysts, programmers, lawyers, designers, researchers, consultants and managers find themselves near the frontier.

Initially they become far more productive.

Then firms discover they need fewer people to produce the same amount.

Then new products absorb some of the freed capacity.

Then the number of tasks AI can perform grows again.

It will probably be an extraordinarily confusing labor market.

And it is possible that economists searching for a clean "AI unemployment event" will miss the deeper transition.

The important quantity may instead be:

How much economically useful output requires a human labor-hour?

That number can fall dramatically long before unemployment reaches extraordinary levels.

10. The 2030s: intelligence crosses into the physical economy

Robotics is likely to enter through the easiest physical environments first.

Factories.

Warehouses.

Ports.

Mines.

Large farms.

Data centers.

Construction prefabrication.

These environments can be partially standardized, making autonomy easier.

The machines do not need to be perfect.

They need to become economically preferable at enough tasks.

And unlike a human employee, a robot's learned policy can potentially improve across the entire fleet.

A mistake observed in Osaka can improve a machine in Ohio.

Physical experience becomes data.

Data improves the model.

The model updates the fleet.

The fleet generates more experience.

Robotics therefore acquires its own data flywheel.

Humanoid form will matter where human infrastructure dominates, but there will be no reason for every machine to resemble us.

Autonomous civilization will contain machines optimized for warehouses, mines, farms, oceans, pipes, skies, operating rooms and orbit.

“Robot” may eventually become as broad and uninformative a category as “computer.”

By the late twentieth century we stopped imagining computers primarily as rectangular boxes performing arithmetic.

Computers dissolved into cars, phones, cameras, factories and appliances.

Physical AI can disappear the same way.

Eventually intelligence is simply a property of machinery.

11. The 2040s and beyond: capital begins building capital

Here lies the true threshold.

Not humanoid robots.

Not AGI.

Not even automated science individually.

The threshold is reached when an automated industrial system can devote a significant fraction of its output to expanding its own productive capacity.

Robots build factories.

Factories make robots.

Robots construct power plants.

Power plants feed compute.

Compute improves the AI coordinating the system.

AI improves the machinery and logistics.

The improved system builds the next generation faster and more cheaply.

This gives us a way to describe the next century without pretending to know its dates.

Progress will migrate from bottleneck to bottleneck.

2026: cognition.

Later: reliability.

Then physical manipulation.

Then manufacturing capacity.

Then electricity.

Then materials.

Then permitting and institutional coordination.

Then perhaps ecological limits.

Each time one bottleneck falls, the next becomes visible.

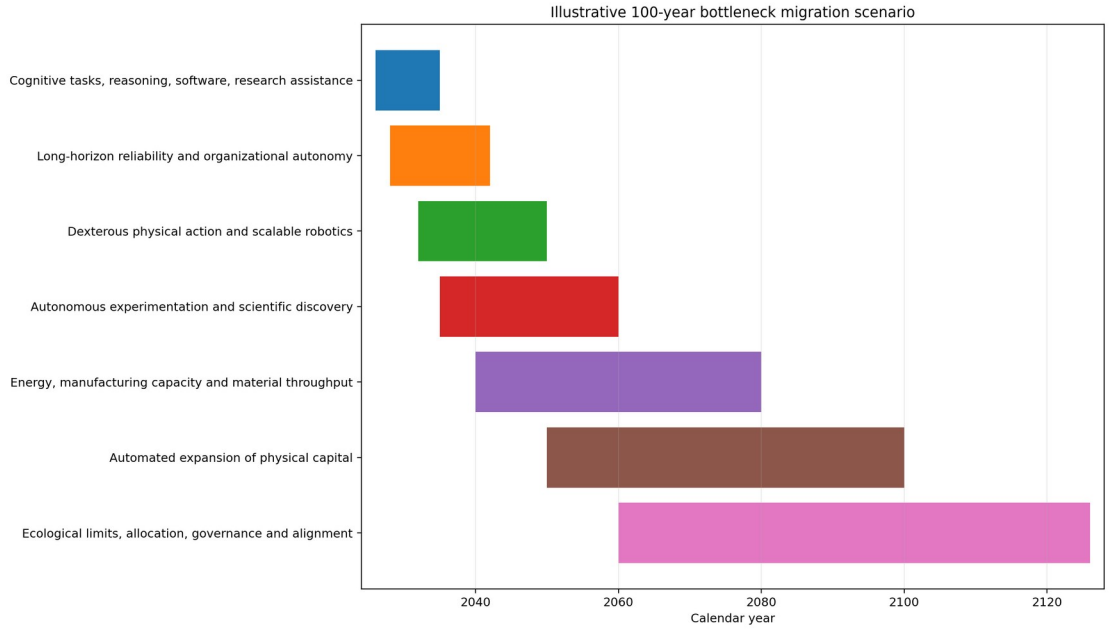

Illustrative bottleneck migration scenario

12. Energy becomes the great physical denominator

This is already happening.

AI feels immaterial because its outputs arrive as words and images.

Its infrastructure is emphatically material.

Data centers consume electricity.

Semiconductor fabrication consumes energy and water.

Cooling consumes energy.

Robots consume energy.

Mining consumes energy.

Manufacturing consumes energy.

Desalination consumes energy.

Recycling consumes energy.

Carbon removal consumes enormous amounts of energy.

Making fuels from atmospheric carbon consumes energy.

Even transforming abundant elements into useful industrial materials ultimately requires energy.

The International Energy Agency's 2026 central outlook puts global data-center electricity consumption around **950 TWh by 2030**, roughly double 2025 levels, while AI-focused facilities grow faster than the sector overall. [P2-14]

This is a perfect example of bottleneck migration.

Intelligence becomes abundant.

Therefore energy becomes relatively scarce.

And what does abundant intelligence do?

It attacks energy scarcity.

Better solar cells.

Better batteries.

Better grid design.

New geothermal techniques.

Advanced fission.

Fusion.

Superconductors.

Power electronics.

Materials.

Control algorithms.

The intelligence-energy feedback loop could become one of the dominant forces of the twenty-first century.

13. Suppose fusion and practical room-temperature superconductivity both arrive

Neither is guaranteed.

Their arrival dates are unknowable.

But asking what happens *conditional on their success* reveals the structure of the system.

A cheap, manufacturable ambient-condition superconductor could dramatically improve electrical transmission, high-field magnets, motors, generators and numerous industrial systems.

Fusion, if it becomes genuinely economical rather than merely scientifically possible, could add another huge source of dispatchable energy.

Together they would not give humanity "free energy."

Thermodynamics remains.

Power plants require capital.

Materials remain finite.

Transmission systems must still be built.

Waste heat still exists.

But the marginal economics of enormous categories of activity could change.

Energy-intensive desalination becomes cheaper.

Synthetic fuels become cheaper.

Recycling becomes more attractive.

Fertilizer becomes cheaper.

Certain forms of industrial heat become easier.

Carbon removal becomes more feasible.

Materials processing becomes cheaper.

Computation becomes cheaper.

Robotic labor becomes cheaper.

And cheaper computation feeds back into the AI systems discovering the next improvements.

The consequence is not merely "electricity bills go down."

The feasible set of civilization expands.

Things previously rejected because they required too much energy become candidates.

14. When knowledge work disappears as a limiting factor

Now confront the uncomfortable case directly.

Suppose AI eventually does essentially all economically useful knowledge work at superhuman quality and negligible marginal cost.

Not merely writing.

Mathematics.

Science.

Engineering.

Medicine.

Law.

Finance.

Management.

Design.

Strategy.

Programming.

Teaching.

Research.

What then remains economically special about the human brain?

Perhaps nothing.

That does not mean human beings become worthless.

Economic usefulness and moral worth are entirely different categories.

Children are not valuable because of their GDP contribution.

Friends are not valued because of labor productivity.

Grandparents do not need a comparative advantage over software to deserve existence.

But modern political economy implicitly connects access to resources with productive labor.

Most people obtain purchasing power primarily by selling their time.

If human labor ceases to be an indispensable productive input, that mechanism becomes unstable.

The central economic question could therefore migrate once again.

From:

How do we produce enough?

To:

Who owns the systems that produce it?

That is a very different civilization.

15. The scarcity inversion

AI and robots cannot eliminate scarcity.

They change **what is scarce**.

If intelligence becomes abundant, intelligence is no longer the principal economic moat.

If labor becomes abundant, labor stops determining prices.

If energy becomes much cheaper, energy-intensive transformations become ordinary.

What remains?

Land.

Rare elements.

Ecological carrying capacity.

Compute hardware.

High-quality physical locations.

Human attention.

Legitimacy.

Trust.

Authenticity.

Political power.

Unique biological ecosystems.

Time.

And things whose value is intrinsically positional.

There can only be one apartment overlooking a particular park.

Only so many people can occupy a pristine beach without destroying its character.

Only one person can own an original artifact.

Only twenty players can simultaneously participate in a particular small-group human experience.

Abundance therefore does not produce a world without economics.

It produces a new economics organized around whatever the machines cannot cheaply replicate.

This happened before.

Industrialization made manufactured goods cheaper and increased the relative value of services.

Computers made computation cheap and increased the value of data and software.

The Internet made information abundant and made attention scarce.

AI makes generated cognition abundant.

Robotics may make executed labor abundant.

Cheap energy may make transformation of matter abundant.

Scarcity keeps moving.

16. 2050–2080: the physical world becomes programmable

At some point the phrase **software eats the world** may acquire a much more literal successor.

Intelligence eats the boundary between software and matter.

Imagine specifying an outcome:

Build 100,000 housing units.

Restore this watershed.

Construct a desalination facility.

Convert this abandoned industrial area into forest.

Build a railway.

Design a vaccine manufacturing facility.

Establish a lunar research station.

Today every one of these statements expands into thousands of subprojects requiring thousands of humans to coordinate.

Future AI systems could recursively decompose such intentions:

survey,

design,

permit,

procure,

manufacture,

transport,

construct,

inspect,

repair,

optimize.

Robotic systems execute the physical steps.

The gap between saying *what we want* and organizing the labor necessary to create it begins collapsing.

The physical world becomes partially **declarative**.

Not:

here are the instructions.

But:

here is the goal.

Find a way.

This is what generative AI did to portions of software creation.

A century of robotics could do it to the built environment.

17. Medicine becomes engineering

The same recursive system changes biology.

Once AI can read essentially the entire scientific literature, reason over genomic and proteomic data, design molecules, model proteins, operate automated laboratories, analyze clinical evidence and continuously update hypotheses, the pace of biomedical experimentation can rise dramatically.

Again, biology will resist.

Humans are not software.

Clinical trials take time.

Organisms are staggeringly complicated.

Evolution produces countless nonlinear interactions.

Safety matters enormously.

But intelligence attacks each layer.

Digital models reduce the candidate space.

Organoids and advanced models test more possibilities.

Robotic labs increase throughput.

Personalized monitoring generates data.

AI designs the next experiments.

Biology gradually becomes more engineerable.

The consequence over a century could range from dramatically better treatment of existing disease to regenerative medicine, organ replacement and substantial changes in healthy lifespan.

The exact achievements are unknowable.

The directional effect is less mysterious:

when the cost of biological hypothesis generation and experimentation falls, more biology becomes technologically accessible.

18. 2080–2126: the remaining frontier may be governance

Forecasting individual technologies a century ahead is mostly fiction.

A person in 1926 could imagine faster airplanes.

They were unlikely to predict CRISPR, smartphones, Wikipedia, GPU clusters and conversational artificial intelligence with useful specificity.

So a responsible forecast of 2126 should not list gadgets.

It should identify constraints.

If advanced AI persists for a century, if civilization survives its associated risks, and if robotic and scientific automation continue improving, it is plausible that many current technical constraints become dramatically less important.

That leaves a class of problems technology cannot solve merely by becoming more capable.

What should be built?

Whose preferences count?

How much wilderness should remain wilderness?

Who controls automated productive capital?

May autonomous systems make political decisions?

What rights, if any, belong to artificial minds?

How should wealth created without human labor be distributed?

How much inequality is acceptable when material deprivation is technically unnecessary?

How should human communities preserve agency when machines can optimize almost everything more efficiently?

How much optimization do we actually want?

These are not engineering problems.

They are problems of values.

And values cannot be derived from greater processing power alone.

A civilization with machines capable of producing almost anything can still choose badly.

It can produce weapons.

Surveillance.

Propaganda.

Ecological destruction.

Extreme concentrations of power.

Optimization toward meaningless metrics.

Abundance amplifies intent.

That means the final bottleneck may be the oldest one.

Wisdom.

19. The dangerous feedback loop

The positive version of the recursive stack is extraordinary:

AI improves science.

Science improves energy.

Energy improves AI.

AI improves robotics.

Robotics increases production.

Production expands science.

But recursion is morally neutral.

The same mechanism can accelerate cyberattacks, autonomous weapons, engineered pathogens, environmental extraction or political control.

An AI that can design a cure may understand biological mechanisms relevant to harm.

A robot capable of construction can also destroy.

Cheap energy can run desalination plants or weapons factories.

Automated persuasion can educate or manipulate.

The speed of the system makes this particularly dangerous because governance historically operates much more slowly than technology.

Parliaments deliberate.

Courts litigate.

Treaties negotiate.

Institutions build legitimacy over decades.

Software updates overnight.

If technological latency approaches months while institutional latency remains years, civilization develops a dangerous phase mismatch.

Part I documented the collapse of **technological diffusion latency**.

Part II therefore reveals another quantity we should begin measuring:

governance latency.

The gap between when humanity acquires a capability and when institutions learn to manage it.

That delta may become more important than the invention-to-adoption delta with which we began.

20. The hundred-year curve

The most useful forecast is therefore not a list of dates.

It is a sequence of bottlenecks.

Approximate era	Dominant bottleneck being attacked
2020s	Cognitive tasks, reasoning, software, research assistance
late 2020s–2030s	Long-horizon reliability and organizational autonomy
2030s–2040s	Dexterous physical action and scalable robotics

Approximate era	Dominant bottleneck being attacked
2030s–2050s	Autonomous experimentation and scientific discovery
2040s onward	Energy, manufacturing capacity and material throughput
mid-century onward	Automated expansion of physical capital
later century	Ecological limits, allocation, governance and alignment

The dates are deliberately fuzzy.

Any one could arrive earlier.

Some may take far longer.

Some may never fully occur.

Room-temperature superconductivity is not guaranteed.

Commercial fusion is not guaranteed.

General humanoid robotics is not guaranteed on a particular schedule.

Total knowledge-work automation is not guaranteed.

But the **direction of bottleneck migration** does not require any one of them.

As intelligence gets cheaper, whatever intelligence cannot yet accomplish becomes relatively more valuable.

AI then gets pointed at that constraint.

If it helps remove the constraint, attention moves to the next one.

This is the continuation of the latency collapse.

21. The first technology that can climb the stack

Every transformative technology in Part I occupied one layer of civilization.

Agriculture transformed food.

Writing transformed memory.

Printing transformed knowledge replication.

Steam transformed mechanical power.

Rail transformed movement.

Electricity transformed energy distribution.

Mass production transformed manufacturing.

Computers transformed calculation.

The Internet transformed communication.

AI begins with cognition.

But cognition is the activity we use to improve all the other layers.

That gives AI an unusual position in the hierarchy.

A better printing press does not independently investigate superconductivity.

A railway does not redesign agriculture.

Electricity does not discover mathematics.

The Internet does not, by itself, synthesize a drug.

AI potentially does all of them.

And once AI controls robots, its outputs no longer terminate at information.

They terminate in **actions**.

Once it controls laboratories, actions produce **discoveries**.

Once it controls manufacturing, discoveries produce **objects**.

Once it controls construction, objects produce **infrastructure**.

Once infrastructure produces more computation and robots, the circle closes.

For the first time, a GPT may be able to climb the entire technological stack.

22. The century in one loop

We can compress the possible next hundred years into one sequence:

Humanity creates machines that can reason.

Those machines help humanity create machines that can act.

Reasoning machines and acting machines automate scientific discovery.

Automated science produces better materials, medicine, energy systems and computers.

Better computers produce more intelligence.

Better energy supplies the machines.

Better robotics constructs the infrastructure.

The infrastructure manufactures more machines.

The machines perform increasing portions of both cognitive and physical production.

Human labor ceases to be the dominant constraint on civilization's productive capacity.

Material scarcity retreats in some domains.

The economy reorganizes around the scarcities that remain.

And eventually the limiting question changes from:

What can we do?

To:

What should we do?

That may be the real endpoint of the latency collapse.

23. The final disappearance of AI

Part I argued that a general-purpose technology reaches cultural maturity when its name begins to disappear.

Nobody boasts that his office uses electricity.

Nobody describes a corporation as computerized.

Increasingly, nobody needs to say that a business is connected to the Internet.

These things became environment.

If the trajectory described here continues, the same thing eventually happens to AI.

The phrase *AI-powered car* disappears because cars are intelligent.

AI doctor disappears because computational intelligence is inseparable from medicine.

AI robot becomes redundant because autonomous machinery necessarily contains intelligence.

AI scientist disappears because science itself consists of networks of biological and artificial investigators.

Eventually asking whether something "uses AI" may sound as peculiar as asking whether a modern hospital uses electricity.

And a child born into that world will not experience any of this as artificial intelligence.

She will experience it as reality.

Her home will contain objects that understand speech.

Infrastructure will respond intelligently to changing conditions.

Machines will maintain machines.

Scientific systems will continuously search possibility spaces no human being could personally comprehend.

Medicine will know her biology in extraordinary detail.

The distinction between digital command and physical execution will feel increasingly thin.

She may look back at our era — the 2020s — with the same mixture of fascination and incomprehension with which we look at the first electrical grids.

She may ask how people possibly organized a city when almost every decision had to pass through a human brain.

How doctors could treat patients while personally reading only a tiny fraction of medical literature.

How scientists could investigate only the experiments their small laboratory could physically perform.

How factories stopped every night because people needed to sleep.

How a society capable of producing plenty nevertheless required almost everyone to sell most of the waking hours of their life merely to obtain the right to consume some of that production.

To her, those arrangements may not look natural.

They may look like technological constraints.

Because they were.

Epilogue: From the latency collapse to the choice explosion

For ten thousand years, human civilization has been primarily constrained by what it could do.

Could we grow enough food?

Could we cross the ocean?

Could we transmit a message?

Could we manufacture enough steel?

Could we compute the trajectory?

Could we discover the molecule?

Could we afford the energy?

Could we find enough skilled people?

Technology repeatedly enlarged the answer.

AI may be unusual because it attacks the mechanism by which we discover new answers.

Robotics attacks the mechanism by which answers become physical.

Autonomous science attacks the mechanism by which ignorance becomes knowledge.

Abundant energy would attack the physical cost of transforming one state of matter into another.

Put those together and something changes at the civilizational level.

The central problem of humanity may cease to be insufficient capability.

We may enter an age of **excess capability**.

A civilization able to attempt vastly more things than previous civilizations could imagine will have to become vastly better at choosing among them.

And this reveals a strange final inversion.

The technological history of humanity begins with scarcity of capability.

It may end — not literally, but conceptually — with scarcity of wisdom.

The last latency to collapse may therefore be one technology cannot collapse for us:

the distance between **what we are capable of doing**

and **what we are wise enough to choose.**

Source Index

Part I

- **P1-01** — Transforming GPTs / Lipsey-Carlaw-Bekar list reproduced by GAO. https://www.gao.gov/assets/gao-14-181sp.pdf
- **P1-02** — Early agriculture and domestication context. https://www.cambridge.org/core/books/worlds-of-the-indian-ocean/birth-of-the-state/B1C0C980A44DFF3EC38C4FFBBA0F7D12
- **P1-03** — Rise of metallurgy in Eurasia. https://www.ucl.ac.uk/social-historical-sciences/archaeology/research/research-directory/material-worlds/rise-metallurgy-eurasia
- **P1-04** — The Origins of Writing. https://www.metmuseum.org/de/essays/the-origins-of-writing
- **P1-05** — Uptake of mechanical technology in the ancient world: watermill. https://oxrep.classics.ox.ac.uk/working_papers/uptake_mechanical_technology_ancient_world_watermill/
- **P1-06** — Cromford Mills / factory system. https://historicengland.org.uk/listing/the-list/list-entry/1488615
- **P1-07** — UK Parliament: rail passenger growth. https://www.parliament.uk/business/publications/research/olympic-britain/transport/plight-train/

- **P1-08** — Smithsonian: Americans adopt the auto. https://americanhistory.si.edu/explore/exhibitions/america-on-the-move/online/americans-adopt-auto
- **P1-09** — Smithsonian: commercial aviation at mid-century. https://airandspace.si.edu/explore/stories/commercial-aviation-mid-century
- **P1-10** — Ford: Highland Park moving assembly line. https://corporate.ford.com/corporate/corporate/articles/history/highland-park/
- **P1-11** — U.S. Census: Home Computers and Internet Use in the United States, 2000. https://www.census.gov/content/dam/Census/library/publications/2001/demo/p23-207.pdf
- **P1-12** — Internet Society: Brief History of the Internet. https://www.internetsociety.org/internet/history-internet/brief-history-internet/
- **P1-13** — Pew: Internet use. https://www.pewresearch.org/chart/internet-use/
- **P1-14** — Stanford HAI: 2026 AI Index, Economy. https://hai.stanford.edu/ai-index/2026-ai-index-report/economy
- **P1-15** — WIPO 2026: How do new technologies diffuse?. https://www.wipo.int/web-publications/world-intellectual-property-report-2026/en/1-how-do-new-technologies-diffuse.html
- **P1-16** — WIPO Working Paper 91 summary: 31 technologies, 139 countries, two centuries. https://www.wipo.int/publications/en/search.jsp?lang=EN&rows=100&set4=184&sort=pubDate
- **P1-17** — WIPO World Intellectual Property Report 2026 executive summary. https://www.wipo.int/web-publications/world-intellectual-property-report-2026-technology-on-the-move-executive-summary/en/executive-summary.html

Part II

- **P2-01** — Nature: AI cracks 80-year-old mathematics challenge. https://www.nature.com/articles/d41586-026-01651-0
- **P2-02** — OpenAI: On the Navier-Stokes Millennium Prize Problem. https://openai.com/index/navier-stokes-solution/
- **P2-03** — OpenAI: Research acceleration - view inside OpenAI. https://openai.com/index/research-acceleration-view-inside-openai/
- **P2-04** — METR: Task-completion time horizons of frontier AI models. https://metr.org/time-horizons/
- **P2-05** — NBER: Chaining Tasks, Redefining Work. https://www.nber.org/papers/w34859
- **P2-06** — IFR: World Robotics 2025 - Industrial Robots executive summary. https://ifr.org/img/worldrobotics/Executive_Summary_WR_2025_Industrial_Robots.pdf
- **P2-07** — Figure: F.02 contributed to production of 30,000 cars at BMW. https://www.figure.ai/news/production-at-bmw
- **P2-08** — NVIDIA Isaac GR00T. https://developer.nvidia.com/isaac/gr00t
- **P2-09** — Nature: An autonomous laboratory for accelerated synthesis of inorganic materials. https://www.nature.com/articles/s41586-023-06734-w

- **P2-10** — npj Computational Materials: AI-accelerated workflow for superconductor discovery. https://www.nature.com/articles/s41524-026-01964-8
- **P2-11** — PPPL: PACMAN AI framework for fusion control. https://www.pppl.gov/news/2026/pacman-ai-framework-controlling-fusion-systems-safely-makes-key-decisions-milliseconds
- **P2-12** — Google DeepMind: Bringing AI to the next generation of fusion energy. https://deepmind.google/blog/bringing-ai-to-the-next-generation-of-fusion-energy/
- **P2-13** — ILO: Workers' exposure to AI. https://www.ilo.org/publications/workers%E2%80%99-exposure-ai-what-indicators-tell-us-%E2%80%93-and-what-they-don%E2%80%99t
- **P2-14** — IEA: Key Questions on Energy and AI. https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary