The Arrow Points the Other Way

JOSH SOLOMON · KNOWN QUANTITY LLC · NOVEMBER 2026

> Everyone asks what happens when AI gets a quantum computer. The evidence so far says that is the wrong question. Quantum computers are not about to make AI faster. AI is already making quantum computers work, and the place where the two really fuse is not the data center. It is the quantum network, where quantum states get moved, stored and measured together. This essay says what is true today, what is likely in a year, and what I think ten, twenty and fifty years look like, with every claim marked as documented, reported, disputed or my own inference.

Words worth knowing

The words this essay leans on, each explained twice: first in plain English, then in one line for the expert.

qubit

The basic unit of a quantum computer, as a bit is for an ordinary one. A bit is always 0 or 1. A qubit can sit in a blend of both, but measuring it gives a plain 0 or 1, so you read out only one ordinary bit per qubit.

> FOR THE EXPERT A two-level quantum system whose state is a normalized superposition of two basis states; describing n qubits takes 2^n complex amplitudes, but by Holevo's bound measuring them yields at most n classical bits.

error correction

Heat, stray signals and tiny flaws keep knocking qubits off course. Error correction spreads one qubit's information across many qubits and keeps checking whether small groups still agree. The checks never read the information itself, since reading would destroy it, but a failed check points to the mistake so it can be fixed.

> FOR THE EXPERT Quantum error correction encodes logical information in an entangled state of many physical qubits and repeatedly measures stabilizer (parity) operators whose outcomes, the syndrome, reveal errors without revealing the encoded state.

decoder

The error checks give clues, not answers, like smoke alarms that say which rooms smell of smoke but not where the fire started. The decoder is the ordinary program that reads the clues and works out the most likely mistake, fast enough to finish before the next round arrives.

> FOR THE EXPERT A classical algorithm (for example minimum-weight perfect matching, union-find, or a trained neural network) that maps a stream of syndrome measurements to an inferred correction under real-time latency constraints.

logical qubit

As the essay puts it, one reliable qubit built out of many unreliable physical ones. Picture a crew rowing one boat: one rower slipping does not stop the boat, because the others make up for it. A bigger crew lets fewer mistakes through, but only if each rower is good enough to begin with.

> FOR THE EXPERT A qubit of information encoded across many physical qubits in an error-correcting code, whose logical error rate falls exponentially with code distance when physical error rates are below threshold; the essay's 'qubits that only detect errors' are error-detecting codes, which flag errors but cannot correct them.

fault-tolerant

A fault-tolerant quantum computer gives right answers even though its parts keep making small mistakes, because errors are caught and fixed faster than they pile up, including mistakes made during the fixing. No full-size one exists yet. Today's error-corrected machines mostly store qubits or run small circuits.

> FOR THE EXPERT A computation is fault-tolerant when it uses encoded gates, measurements and syndrome extraction designed so that a single physical fault cannot spread into an uncorrectable logical error, allowing arbitrarily long computations once physical error rates are below threshold.

density functional theory

A method computers use to predict how the electrons in a molecule or material behave, which tells chemists how it holds together and reacts. The exact problem would take far too long, so it uses an approximation that usually works but fails on some substances, such as metal centers in enzymes.

> FOR THE EXPERT A method that computes ground-state properties from the electron density instead of the many-electron wavefunction; exact in principle, it relies in practice on approximate exchange-correlation functionals that fail for strongly correlated systems such as many transition-metal complexes.

entangled

Two quantum things are entangled when they behave like one linked pair, so measuring one tells you something about what you will get from the other, even far apart. It is not like gloves mailed in two boxes, because the results are not fixed in advance, and it cannot send a message faster than light.

> FOR THE EXPERT A state of two or more systems is entangled when it cannot be written as a product, or a mixture of products, of states of the parts; its correlations can violate Bell inequalities yet permit no signalling.

quantum memory

A device that holds a quantum state, such as the state of one particle of light, and hands it back later without reading it, since reading would change it. Today's memories hold states for tiny fractions of a second and often lose them; the Geneva one returned about one photon in five.

> FOR THE EXPERT A system (for example a rare-earth-doped crystal, trapped atoms or a diamond color center) that stores a quantum state on demand and retrieves it later, judged by efficiency, storage time, fidelity and number of modes.

quantum network

Ordinary networks move bits. A quantum network moves quantum states, usually carried by single particles of light in fiber, between distant machines without reading them. Light fades in long fiber, and a quantum state cannot be copied and boosted like a phone signal, so long networks need relay stations called repeaters. None yet beats plain fiber.

> FOR THE EXPERT A network that distributes entanglement or transmits quantum states between nodes such as processors, memories and sensors over optical channels; because of loss and the no-cloning theorem, long distances require quantum repeaters built on entanglement swapping and, in advanced designs, error correction.

RSA-2048

A way of locking information that rests on one fact: multiplying two huge prime numbers is easy, but working back from the result to the two primes would take ordinary computers far longer than anyone can wait. 2048 is the key size in bits. A big enough fault-tolerant quantum computer could work backward.

> FOR THE EXPERT The RSA public-key cryptosystem with a 2048-bit modulus N = pq, whose security rests on the presumed hardness of factoring N classically; Shor's algorithm factors it in polynomial time on a fault-tolerant quantum computer.

Where I am standing

I am not a physicist. I run a small company where an AI does most of the logistics. So I came to this the way most people do, assuming the story was simple: quantum computers are absurdly fast, AI needs absurd amounts of computing, put them together.

Then I read the papers, including the ones written by the people who build the machines. The simple story does not survive them. A stranger and better one does.

The merge everyone is waiting for is not coming

Start with the dream: train the next giant AI model on a quantum computer.

The strongest argument against it comes from inside the field. In 2023 Torsten Hoefler, Thomas Häner and Matthias Troyer, who works on quantum computing at Microsoft, compared a generous future quantum computer with a single ordinary graphics chip. The quantum machine moves data in and out roughly ten thousand times slower, and does basic arithmetic billions of times slower. It only wins when the algorithm itself gives an enormous speedup on a problem that needs very little data. Their verdict on big-data problems and on most machine learning training was blunt: unpromising. (Documented.)

Google's own quantum team reached the matching conclusion two years earlier. A quadratic speedup, the kind you get for general search and optimization, is eaten by the cost of error correction. You need something far steeper. (Documented.)

The rest of the wall is built from three more bricks. A qubit register needs an exponential number of numbers to describe, but you can only read out one ordinary bit per qubit, so "qubits store exponential data" is false in the way that matters. Loading a large classical dataset into a quantum state costs about as much as just processing it. And when quantum machine learning algorithms did look exponentially faster, Ewin Tang showed in 2018 that a clever classical algorithm could match one of the flagship examples, and others extended that widely. (All documented, and old enough to be textbook.)

So my first claim, as plainly as I can put it: nothing published gives a credible path to training or running today's kind of AI model on a quantum computer. Not next year and probably not in fifty. (My inference, from documented parts.)

The merge that already happened

Now turn the question around. What is AI doing for quantum computers?

A quantum computer is a chip full of fragile qubits wrapped in a much larger classical machine. It checks for errors in rounds, about a million rounds a second on superconducting chips and far fewer on trapped ions and neutral atoms, works out what went wrong, and keeps track of the fixes in software. That wrapper is where the merge has already happened.

Three examples from the last twelve months.

Decoding. Error correction produces a stream of clues about which qubits went wrong. Something has to read those clues faster than the next round arrives. Google DeepMind's first neural decoder, in 2024, was more accurate than the standard methods but too slow to keep up. The follow-up preprint in December 2025 reports a decoder fast enough to keep pace with a new round every millionth of a second, timed on Google's own AI chips, for codes bigger than any yet run on hardware. On real data recorded from Willow it was as accurate as the slow first version. It has not been wired into a running chip yet, and each answer still takes hundreds of millionths of a second to come back. (Documented from the paper.)

Calibration. In a paper reported in Nature this July, a reinforcement learning agent adjusted more than a thousand control settings on Google's Willow chip while it ran error correction. The chip had already been through Google's full standard calibration. The agent still cut the logical error rate by a further 20 percent. These were repeated runs of an error-corrected memory, not one long computation. (Documented from the preprint.)

Running the lab. In September, a group at MIT working with OpenAI had a general-purpose AI agent bring up a simple six-qubit chip nobody had measured before, on real hardware. It found the resonators, ran the standard measurements, and reported the numbers. On its four simpler qubits, humans stepped in on 4 of 40 measurements. (Documented.) The same report is honest about the limits: on harder devices it needed real guidance, accepted one flawed scan, and was slower than an experienced researcher.

Then there is design. A preprint this summer describes an AI system searching for new error-correcting codes. In simulation, its best codes beat the reference codes IBM is building its roadmap on. (Documented as a preprint, simulation only, not peer reviewed, and nobody has run one on a chip.)

The honest objection matters here. Learned control is not proven to be necessary. IBM decodes its codes in under half a millionth of a second with no machine learning at all. So I will not say a quantum computer needs AI to run. I will say this: the best-performing control stacks we have already contain it, and every clear arrow between these two technologies today points from AI to quantum, not the other way. (Inference.)

Who gets replaced first

There is an uncomfortable reading of everything above, and I would rather say it than leave it implied.

For a long time, the hard part of this field was assumed to be safe from automation. Tuning a quantum device was a craft. It took years in a lab to learn which knob to turn when a measurement looked wrong.

Look at what the last year did to that assumption. An AI agent brought up a chip nobody had measured before. (Documented.) A second one, at the University of Chicago, read a published physics paper and reproduced its benchmark measurement on real hardware in about three minutes. (Documented.) A third calibrated a 20-qubit commercial chip by reading the tuning plots the way a graduate student would. (Reported.) A learning agent improved on a chip that had already been through Google's full calibration. A search program proposed error-correcting codes that beat, on paper, the ones a company roadmap is built on. And on the theory side, a learning system matched the best circuits humans had constructed by hand. (Documented.)

None of these systems understands quantum mechanics better than a physicist. They do not need to. The work they are taking is the procedural middle of the job: the bring-up, the calibration, the search through a space too big to search by hand. That is most of the hours.

The honest limits are in the same reports. The agents were slower than an experienced researcher on the harder devices. One accepted a bad measurement without noticing. They needed a human to widen a search range. So the thing that is still scarce is not the knowledge of how to run the experiment. It is the judgment to know when the machine running it is wrong.

I do not say this from a safe seat. I wrote an essay this week about handing my own logistics to an AI, and the tedious work I used to do by hand was the first thing it took. The pattern is the same in a physics lab as in a finance department. The procedure goes first, however hard it was to learn. What is left is taste, judgment, and the choice of what is worth doing. Nobody's expertise is exempt from that, and the people who see it early will be the ones still steering.

What quantum will actually give AI

If not speed, then what? I think two things, and neither is what the headlines promise.

The first is ground truth. AI models for chemistry and materials are trained on answers computed by density functional theory, the workhorse method chemists use because it is fast, and which gets its speed from an approximation. Meta's 2025 dataset for this contains more than 100 million such calculations and took over six billion hours of computing. (Reported.) The models built on it are extraordinary, and they inherit the approximation's blind spots. The cases where it fails are the ones chemists care most about: metal centers in enzymes, catalysts, some battery materials.

Those are the cases where a fault-tolerant quantum computer could in principle compute the energy as accurately as you are willing to pay for, within a chosen model of the molecule, as long as you can hand it a reasonable starting guess. So the realistic picture is not a quantum computer replacing the AI model. It is a quantum computer producing a small number of answers nobody else can produce, and a classical AI model learning from them and then running a billion times. The quantum machine is the oracle. The AI is the student. (Inference.)

The objection is real. Each of those answers still costs billions of logical operations and days of machine time on a computer that does not exist yet. And a careful 2023 review found no evidence so far that quantum computers have an exponential edge across chemistry in general. Classical methods, including AI ones, keep getting better. There is a race here, and the classical side is not standing still. (Documented.)

The second gift is stranger, and it is the one I would most like to put to the people who build this hardware.

The part that happens on the network

There is a small set of results where quantum machines provably learn better than classical ones. They share a feature. The data is quantum to begin with.

In 2022 a team including Hsin-Yuan Huang, John Preskill and Jarrod McClean showed that for certain questions about an unknown quantum state, a learner that can hold two copies at once and measure them together needs a handful of experiments, while any learner that measures one copy at a time needs a number that doubles with every qubit in the state. They showed it on Google's chip with test states of up to 20 qubits plus 20 qubits of memory, all prepared on the same chip, so nothing came in from outside. Their own paper pictures the next step as moving quantum data from real physical systems straight into a quantum memory. (Documented.) A photonic experiment reported in Science in 2025 did a version of it with light: by keeping an entangled partner beam and measuring the two together, it learned a 100-channel noise process with about a trillion times fewer samples. (Documented from the abstract.)

Read that again as an engineer. The advantage does not come from a faster processor. It comes from measuring copies together, or against an entangled partner you kept, instead of measuring each one alone. That takes a way to move a quantum state from the thing you are studying into a memory without scrambling it, and a processor to do the joint measurement. And here is the catch I did not expect: the advantage grows with the number of qubits you are learning about, but so does the damage from moving them. A lost photon you know you lost only slows you down. A quiet error on each qubit adds up. So doing this at scale needs the same thing a quantum repeater needs, which is moving quantum states with errors you can detect and eventually correct. When the thing you are studying is far from the processor, that is a quantum network.

So here is the idea I would defend. The place where AI and quantum truly merge is that chain: the quantum thing being studied, a clean way to move its state, a memory, a small processor for the joint measurement, and a learned system doing all the classical work around it, the calibration, the drift tracking, the scheduling, the readout and the model the measurement feeds. Quantum does not give AI a faster brain. It may give it a new sense organ: for some questions about quantum things, like the output of a quantum simulator or a strange material, a handful of joint measurements can replace a number of ordinary ones that doubles with every qubit. For a sensor reading an ordinary field, like a magnetometer, the gain is much smaller and nothing like exponential, and the first real uses will probably be quantum

devices checking each other and distant sensors linked together, like telescopes. This is not my idea alone. It is how the people who proved these results describe where they lead. What I would add is that it makes the network people's hardest problem, moving quantum states cleanly, the same problem as the learning one. (Inference, built on documented theory.)

The gap between that sentence and a working device is enormous, and people in the field know the numbers better than I do. A Geneva group this year stored one photon of an entangled pair in a crystal memory, in more than eight thousand time slots at once, while its partner traveled 5.66 kilometers of city fiber. About 19 percent of stored photons came back out, after 63 millionths of a second. (Documented.) Every network result I found is one or two links. Nobody has yet built a chain of repeaters that beats sending the light straight down the fiber at a useful rate. In the first algorithm run across two linked quantum modules, at Oxford in 2025, the gate between the modules was 86 percent accurate, against better than 99.9 percent for gates inside one module. (Documented from the abstract.) The closest anyone has come to a real signal feeding quantum memories is a 2025 Harvard experiment from Mikhail Lukin's group: two diamond memory stations, entangled and up to 1.55 kilometers of fiber apart, caught weak light and compared its phase better than either could alone, the first step toward telescopes linked by a quantum network. Separately, a Cornell and MIT team this summer made a single qubit both sense a magnetic field and process it before reading it out. Nobody has yet moved a sensor's quantum state into a separate processor and learned from it with a proven advantage. (Documented.)

One more result I cannot leave out, because it cuts against my own tidy story. In April this year, Google's theorists and Preskill posted a paper claiming an exponential quantum advantage on plain classical data, streamed past a small quantum computer. It is an advantage in memory, not speed, and their examples use fewer than 60 logical qubits. (Documented as a preprint; not yet independently tested.) If it holds, "quantum only helps with quantum data" is too strong. I am flagging it and not building on it.

The timeline, honestly

Today. Error correction has passed its first real test: making the code bigger makes the stored qubit's error rate drop, by about 2x per step on Google's chip and on Harvard's neutral atoms, and reportedly 12x per step on Quantinuum's ions. (Documented, the ion number reported.) A logical qubit is one reliable qubit built out of many unreliable physical ones, and that is the only kind that counts for real work. These are mostly memories and small circuits, not long computations. Headlines about "94 logical qubits" describe qubits that only detect errors, with a quarter or fewer of the runs kept for the headline result, not fully corrected ones. (Reported, and disputed in framing.) Microsoft's topological qubit is still disputed by named physicists. No quantum computer has done anything commercially useful. AI is in the decoders, the calibration and the code search.

The next few months. IBM has said it expects a quantum advantage that outside groups can verify by the end of this year. I could not confirm whether that or its 2026 chip has landed. Treat it as a test of whether roadmaps mean anything.

One year. The published roadmaps say IBM links modules, Quantinuum ships its next machine, IonQ promises ten thousand qubits on a chip. My guess: the first circuits with tens of fully corrected logical qubits doing thousands of operations, and AI agents calibrating chips as a routine. Scott Aaronson wrote last year that a fault-tolerant machine running Shor's algorithm before the 2028 presidential election was "a live possibility," and clarified that even factoring 15 would count, as long as it was fully fault-tolerant. (Documented.) That is the right size of expectation.

Ten years. Here the numbers get serious. In 2019 the best estimate for breaking RSA-2048, the standard behind many web certificates, was 20 million noisy qubits running for eight hours. In 2025 Craig Gidney at Google cut it to under one million running for under a week. One preprint this year claims under a hundred thousand, and another claims about 26,000 atoms for the elliptic-curve keys most web connections use today, which are cheaper to break than RSA. (Documented, with the newest ones unreviewed and resting on hardware nobody has built.) A survey of 26 experts published this March put the chance of a machine breaking RSA-2048 in under a day within ten years at 28 to 49 percent, the highest in the survey's history. (Reported.) The US government stops allowing the old encryption after 2035. Traffic recorded today whose key exchange is not yet post-quantum can be decrypted later, and that can include AI model weights moving between data centers. If any one of the 2029 roadmaps lands within a few years of schedule, 2036 has machines that can do the first real chemistry, and AI agents write most of their programs because no human can lay out a computation that size by hand. (Inference.)

Twenty years. No company roadmap goes this far. If fault tolerance works, I expect the main job of quantum computers to be simulating matter and feeding the answers to classical AI, with quantum networks linking the machines and the sensors. Training AI stays classical. (Inference.)

Fifty years. Some things do not change at any horizon, because they are theorems or close to it. Quantum computers do not try every answer at once. They do not solve every hard problem. They do not make anything conscious. What I think is plausible is that quantum processors become ordinary attachments, like graphics chips are now, each one run by a learned controller, with the most valuable ones sitting at the end of a sensor. (Inference, and I hold it loosely.)

Why it could all be slower. Long error-corrected computations, as opposed to memories, have never been shown at the error rates useful work needs. Every recent drop in the estimates assumes codes and connections not yet built at scale. Rare bursts of correlated errors may set a floor. Gil Kalai has argued for years that noise will defeat the whole project, and the current results are not yet strong enough to prove him wrong. Fidelity targets have a history of slipping.

Why it could be faster. Error suppression is exponential and now shown three different ways. The requirements keep falling faster than the hardware grows, so the two curves are closing from both ends. And the thing shortening the calibration, the decoding and the code design is AI, which is improving on its own schedule.

That last point is the whole essay in one line. AI is speeding up quantum computing far more than quantum computing will ever speed up AI.

What I would ask a physicist

1. The two-copy learning advantage grows with the number of qubits you learn about, but so does the error in moving them. If each qubit arrives 99 percent intact, a 50-qubit state arrives whole only about 60 percent of the time. Does that mean quantum-enhanced learning at any interesting size needs error-corrected transfer, the same thing the most advanced repeater designs need?
2. In your hardware, which errors announce themselves, so you know a try failed and just try again, and which are silent? Which silent ones would poison a joint measurement without you noticing?
3. Crystal memories hold thousands of time slots at about 20 percent efficiency for tens of millionths of a second. Diamond nodes hold one or two qubits with a much better connection to light and a far longer memory. Which would you build a learning node from, and why is that a different answer from a repeater node?
4. Harvard's 2025 experiment compared weak light caught at two entangled diamond stations 1.55 kilometers apart. What would it take to beat an ordinary telescope array on a real star, and how many years away is that?
5. Where is machine learning already inside your experiment's loop? Which job would you refuse to hand to a learned controller because you could not tell when it was wrong, and what test would change your mind?

Sources and a note on method

Grades used above: **documented** means I or my research assistant opened the paper or its abstract; **reported** means a secondary summary was read, not the paper; **disputed** means named experts disagree; **inference** is my own reasoning.

Against the dream: Hoefler, Häner and Troyer, "Disentangling hype from practicality," Communications of the ACM (2023), arXiv:2307.00523. Babbush and colleagues, "Focus beyond quadratic speedups for error-corrected quantum advantage," PRX Quantum (2021), arXiv:2011.04149. Tang, "A quantum-inspired classical algorithm for recommendation systems" (2018), arXiv:1807.04271. Cerezo and colleagues, Nature Communications (2025), arXiv:2312.09121. The limit on reading out qubits is Holevo's bound.

AI inside the machine: Senior and colleagues (Google DeepMind and Google Quantum AI), "A scalable and real-time neural decoder for topological quantum codes," arXiv:2512.07737. Sivak, Morvan and colleagues, reinforcement learning control on Willow, Nature (July 2026), arXiv:2511.08493, read in summary. Yankelevich and colleagues (MIT and OpenAI), "Agentic Calibration of Superconducting

Qubits" (September 2026). OmniQEC, arXiv:2607.25865. IBM's sub-480-nanosecond decoding as reported in November 2025. A University of Chicago-led group, "Large Language Model-Assisted Superconducting Qubit Experiments," arXiv:2603.08801. IQM and NVIDIA's automated 20-qubit calibration, as reported in October 2026. AlphaTensor-Quantum, Nature Machine Intelligence (2025).

What quantum gives back: Meta's OMol25 dataset, arXiv:2505.08762. Lee and colleagues, "Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry," Nature Communications (2023). Huang and colleagues, "Quantum advantage in learning from experiments," Science (2022), arXiv:2112.00778. Zhao and colleagues, "Exponential quantum advantage in processing massive classical data," arXiv:2604.07639.

Networks and sensing: Geneva multimode memory, arXiv:2608.13177. Main and colleagues (Oxford), distributed quantum computing, Nature (2025). Stas and colleagues (Harvard), "Entanglement assisted non-local optical interferometry in a quantum network," arXiv:2509.09464. Sen and colleagues (Cornell and MIT), "Quantum sensors that compute," arXiv:2608.17400. Liu and colleagues, "Quantum learning advantage on a scalable photonic platform," Science (2025), arXiv:2502.07770. Cain and colleagues on neutral-atom factoring, arXiv:2603.28627.

Timeline: Google's Willow below-threshold result, Nature (2024), arXiv:2408.13687. Quantinuum's 94 and 48 logical qubit result, arXiv:2602.22211. Gidney, "How to factor 2048 bit RSA integers with less than a million noisy qubits," arXiv:2505.15917. The Pinnacle architecture, arXiv:2602.11457. Global Risk Institute Quantum Threat Timeline Report (March 2026), read in summary. NIST IR 8547. Aaronson's blog post of November 13, 2025. Preskill, "Beyond NISQ: the megaquop machine," arXiv:2502.17368.

What I could not check: whether IBM delivered its 2026 chip or its verified advantage; whether the Willow reinforcement learning paper's Nature version matches the preprint I relied on; Quantinuum's ion numbers beyond a press report; and the Oxford and Global Risk Institute sources, which I read only in summary. If any of those are wrong, the correction goes at the top of this page.

I used AI to research and draft this, under my direction, with a standing instruction to tell me where I was wrong before telling me where I was right. The first draft was handed to a reviewer told to read it as a quantum-networking physicist and embarrass me. It found eight factual errors and five sentences of loose physics, and it pointed out that my central idea was already in the papers I was citing. All of that is corrected here, and the idea is now credited to the people who had it first.