Research & briefings 2026

Twelve cited deep dives on DARPA programs, autonomous AI discovery and HAYAT, photonic computing, quantum security, quantum AI, brain-computer interfaces, CRISPR and connectomics: what each is, how it works, where it stands, how it could change life, scenarios to 2040, risks, FAQs and glossaries.

Mohsin Khan Shaik
AI Infrastructure & Site Reliability Engineer · San Francisco · mohsinkhanshaik.com
Edition of 2026-11-22

Contents

01 Computers made of living neurons: DARPA's O-CIRCUIT — DARPA · 6 min

02 Autonomous hypothesis, research and discovery — from DARPA to HAYAT — AI & discovery · 5 min

03 Photonic computers: doing AI math with light — Computing · 4 min

04 Quantum cybersecurity: the migration that can't wait — Quantum · 5 min

05 Atomically precise: DARPA APM and the road from atoms to products — DARPA · 4 min

06 DARPA DICE: 100,000 AI agents and no boss — DARPA · 4 min

07 How DARPA research changes the world — and what's next — DARPA · 4 min

08 DARPA Firefox: reading the brain with light, no surgery — Brain & neurotech · 3 min

09 Neuralink and the brain-machine interface race — Brain & neurotech · 4 min

10 CRISPR and intelligence enhancement: science, limits and ethics — Biotech · 4 min

11 Quantum AI: what's real, what's next — Quantum · 3 min

12 Princeton's 3D maps of the mind: FlyWire and MICrONS — Brain & neurotech · 4 min

Facts are cited to primary or reputable sources listed under each briefing. Timelines, scenarios and "My take" sections are the author's analysis and opinion, not predictions of record. Program details reflect public information available on 2026-11-22.

01 · DARPA

Computers made of living neurons: DARPA's O-CIRCUIT

A deep dive into organoid computing: how living neurons compute, what DARPA's O-CIRCUIT program asks for, who else is building 'wetware', and how it could change AI, medicine and defense.

42 months	~20 W	Mar 2025	800,000
planned program length	power of a human brain	first commercial neuron computer (CL1)	neurons in DishBrain that learned Pong

Key takeaways

- O-CIRCUIT is a 2026 DARPA program to build computers whose computing medium is living neurons and organoid tissue, aimed at low-power AI at the edge.
- It is not trying to build a brain: DARPA says so explicitly, and purely in-silico (silicon-only) proposals are out of scope.
- The field is real but young: Cortical Labs' DishBrain learned Pong in 2022, its CL1 became the first commercial neuron-based computer in 2025, and FinalSpark rents organoid time over the internet.
- The biggest obstacles are engineering, not ideas: keeping tissue alive, reproducible and observable for months.
- Ethics will matter early: as organoids become more complex, questions about sentience and donor consent need clear rules.

Background: why compute with biology?

Modern AI is limited by energy. Training and running large models takes data centers measured in hundreds of megawatts, while the human brain does far more flexible learning on roughly 20 watts. Biological neurons compute with chemistry and electrical spikes, learn continuously, and repair themselves, properties silicon struggles to match.

The term **organoid intelligence** was coined by a Johns Hopkins group including Thomas Hartung to describe using three-dimensional cultures of human brain cells, called *brain organoids*, as a computing substrate [5]. Organoids are grown from stem cells and self-organize into tissue with some of the cell types and layered structure of a developing brain, though they are tiny (millimeters) and far simpler than a real brain.

What O-CIRCUIT asks for

O-CIRCUIT stands for *Organoid Cytomorphic Intelligence Resulting from Convergent Understanding and Information Transfer*. DARPA's solicitation (DARPA-PS-26-28) was published on 1 May 2026 after a proposers workshop on 10 April 2026 [1][2]. Key points:

- **Goal:** power-efficient computers that use living neurons or organoid tissue for AI processing at the tactical edge, where size, weight and power are tight [1][3].

- **Structure:** a 42-month effort with an 18-month first phase, then two 12-month phases, organized around two development approaches, *architecture* and *action* [1].
- **Boundaries:** the FAQ states the program "does not seek to build a brain", and purely in-silico solutions are out of scope [2].
- **Example use:** public coverage describes drones that could "smell" danger such as explosives or chemical agents using olfactory receptor neurons [4].

How organoid computing works

An organoid computer has three layers:

- **The tissue:** neurons (2D cultures or 3D organoids) grown on or around a **multi-electrode array (MEA)**, a chip with many tiny electrodes that can both record spikes and deliver electrical stimulation.
- **The interface:** software encodes inputs as patterns of stimulation (for example, where a ball is in a game) and decodes the network's spiking activity as outputs.
- **The learning loop:** the tissue adapts through synaptic plasticity. In DishBrain, predictable feedback after a "good" action and unpredictable noise after a "bad" one was enough for the culture to improve at Pong [6].

A related approach, **reservoir computing**, treats the living network as a rich, nonlinear "reservoir": you don't train the tissue directly, you train a simple readout on top of its responses. Indiana University's 2023 *Brainoware* used this idea for simple speech tasks, reporting about 78% accuracy at identifying speakers in a narrow test [7].

State of the art in 2026

- **DishBrain (2022):** about 800,000 human and mouse neurons on a CMOS chip learned to play Pong, the field's best-known demonstration [6].
- **Cortical Labs CL1 (March 2025):** launched at Mobile World Congress as the first commercially available computer based on human neurons grown from stem cells, priced at about US$35,000, with remote "wetware-as-a-service" access [6][8].
- **FinalSpark:** a Swiss startup that keeps organoids alive for up to six months and lets researchers run experiments remotely, with a stated goal of AI at a tiny fraction of today's energy [5].
- **DARPA O-CIRCUIT (2026):** the first large defense program to treat organoid tissue as the computing medium itself [1].

Honest caveat: these systems solve toy problems today. Learning Pong is a long way from complex decision-making, and no organoid system rivals a modern neural network on real tasks [7].

How it could change life

- **Edge AI and sensing:** sensors that classify smells, chemicals or signals locally on milliwatts, useful for disaster response, food safety and environmental monitoring.
- **Medicine:** the same platforms are disease models. Organoids from a patient's own cells could show how their neurons respond to a drug, reducing trial and error for epilepsy or neurodegenerative disease.
- **Fewer animal tests:** human neural tissue on a chip can replace some animal experiments in neuroscience and toxicology.

- **Energy:** if hybrid bio-silicon accelerators work for niche tasks, they could cut the energy cost of always-on intelligence.

Future scenarios

2027–2028	2030	2035	2040+
O-CIRCUIT Phase 1 prototypes: tissue that reliably learns simple tasks and talks to electronics; better life-support and longevity.	Field demonstrations of organoid sensing (smell, signal classification) on mobile platforms; standard benchmarks emerge.	Hybrid bio-silicon accelerators for ultra-low-power niches, if reproducibility and lifetime are solved.	Mature biocomputing used alongside silicon for adaptive, always-learning tasks; strong ethical governance in place.

- **Optimistic:** organoid co-processors become a new device class for edge sensing and continual learning.
- **Base case:** biocomputing stays a research and drug-discovery tool, with a few specialized sensing products.
- **Pessimistic:** lifetime, variability and ethics stall the field; neuromorphic silicon captures most of the benefits instead.

Effects and risks

- **Engineering:** tissue needs nutrients, gas exchange, sterile conditions and stable temperature; it ages and varies from batch to batch.
- **Observability:** debugging a living network is far harder than reading a stack trace; new tooling is needed.
- **Ethics:** donor consent, potential for sentience in complex organoids, and how such systems may be used. Clear thresholds and oversight boards will be needed.
- **Dual use:** early applications are military; transparency helps public trust.

What to watch

- O-CIRCUIT award announcements and Phase 1 results.
- Organoid lifetime records (months to years) and batch-to-batch variability.
- Independent benchmarks comparing biological and neuromorphic chips on energy per task.
- Ethics frameworks from bodies such as national academies on organoid research.

Connection to my work

My day job is running AI infrastructure reliably: capacity, power, cooling, observability and incident response. Biocomputing turns every one of those into a biology problem. A "living server" needs health checks (is the tissue alive and responsive?), drift detection (is it still behaving like yesterday?), and graceful degradation when a culture dies. I'd love to see the field borrow SRE practices: SLOs for tissue health, canary cultures, and runbooks.

My take

O-CIRCUIT is a smart bet because it targets a niche where biology has a real edge: low-power, adaptive sensing. I don't expect organoids to replace GPUs. I do expect them to teach us how to build more efficient AI, and to produce useful medical tools along the way. The winners will be the teams that make living computers boringly reliable.

Key players & programs

DARPA O-CIRCUIT	2026 program; organoid tissue as the computing medium for edge AI [1]
Cortical Labs (Australia)	DishBrain (Pong, 2022); CL1 commercial neuron computer (2025) [6][8]
FinalSpark (Switzerland)	Remote-access organoid platform; organoids kept alive up to six months [5]
Indiana University	Brainoware reservoir computing with organoids (Nature Electronics, 2023) [7]
Johns Hopkins	Coined 'organoid intelligence' and set out the research agenda [5]

Frequently asked questions

What is DARPA O-CIRCUIT?
A 2026 DARPA program to build computers that use living neurons and organoid tissue as the computing medium, aimed at power-efficient AI at the tactical edge. It runs for a planned 42 months in three phases.

Is O-CIRCUIT trying to build an artificial brain?
No. DARPA's FAQ states the program does not seek to build a brain. It focuses on harnessing how biological neurons process information efficiently.

Has a biological computer ever learned a task?
Yes. Cortical Labs' DishBrain, about 800,000 neurons on a chip, learned to play Pong in 2022, and its CL1 became the first commercial neuron-based computer in 2025.

Why use living neurons instead of chips?
Energy and adaptability. A human brain runs on about 20 watts and learns continuously, while large AI systems need data centers.

Glossary

Organoid: A small, self-organizing 3D culture of cells grown from stem cells that mimics some features of an organ.

Multi-electrode array (MEA): A chip with many electrodes that records neuron activity and delivers stimulation.

Reservoir computing: Using a complex dynamic system as a fixed 'reservoir' and training only a simple readout layer.

Edge AI: Running AI on devices in the field rather than in a data center.

References

1. DARPA — O-CIRCUIT program page
   https://www.darpa.mil/research/programs/o-circuit
2. DARPA — O-CIRCUIT FAQ (May 2026)
   https://www.darpa.mil/sites/default/files/attachment/2026-05/faq-darpa-program-o-circuit.pdf
3. Military Aerospace — Insect brain-level computing for low-power edge AI
   https://www.militaryaerospace.com/computers/article/55361824/researchers-eye-insect-brain-level-computing-to-bring-low-power-artificial-intelligence-ai-to-the-edge
4. The Sociable — O-CIRCUIT: drones that smell danger
   https://sociable.co/military-technology/darpa-o-circuit-program-drones-smell-danger-biologically-inspired-computer/amp/
5. National Geographic — Scientists want to build 'living' computers powered by brain cells
   https://www.nationalgeographic.com/science/article/brain-cells-organoids-computers-ai-energy

6. IEEE Spectrum — A biological computer for sale (Cortical Labs CL1)
   https://spectrum.ieee.org/biological-computer-for-sale
7. Scientific American — These living computers are made from human neurons
   https://scientificamerican.com/article/these-living-computers-are-made-from-human-neurons
8. Cosmos — First biological computer made from neurons
   https://cosmosmagazine.com/technology/computing/first-biological-computer-neurons
9. arXiv — Brain organoid computing: an overview (2025)
   https://arxiv.org/html/2503.19770v2

02 · AI & DISCOVERY

Autonomous hypothesis, research and discovery — from DARPA to HAYAT

A deep dive into AI that does science: from the first robot scientist to the 2026 'AI scientist' papers, DARPA's expMath and MATHBAC, AlphaEvolve, and how I build verifiable research into HAYAT.

2026	48	expMath	7,865
AI-scientist papers in Nature	multiplications: AlphaEvolve's 4×4 complex result	DARPA's AI co-author for math	tests behind HAYAT's agent loop

Key takeaways

- AI now takes part in every step of research: reading literature, proposing hypotheses, planning experiments, and writing papers.
- 2026 was a milestone year: Nature published papers on Google's AI co-scientist, FutureHouse's Robin and Sakana's AI Scientist, while critics note humans still did the lab work and much output lacks novelty.
- DARPA is funding AI co-authors for mathematics (expMath) and self-improving agentic discovery platforms (MATHBAC).
- The bottleneck is shifting from generating ideas to verifying them: provenance, citations and honest uncertainty.
- HAYAT applies the same principles locally: cited research, background learning, action verification and 7,865 automated tests.

A short history of machines doing science

- **1970s–80s:** expert systems such as DENDRAL inferred molecular structures from mass-spectrometry data.
- **2009:** the robot scientist "Adam" (Aberystwyth University) autonomously generated and tested hypotheses about yeast genes.
- **2020–2024:** AlphaFold predicted protein structures at near-experimental accuracy, recognized by the 2024 Nobel Prize in Chemistry.
- **2025:** DeepMind's AlphaEvolve, an LLM-driven evolutionary coding agent, found a way to multiply 4×4 complex-valued matrices with 48 scalar multiplications, improving on Strassen's 1969 approach for that case, and sped up a Gemini kernel by 23% [6].
- **2026:** "AI scientist" systems are published in Nature [1].

How an AI scientist works

Most systems are **multi-agent loops** built on large language models:

1. **Literature agent:** searches and summarizes papers, extracting claims and evidence.

2. **Hypothesis agent:** proposes ideas, often many in parallel, and ranks them through debate or tournament-style comparison.
3. **Planning agent:** designs experiments or analyses; in self-driving labs this becomes robot instructions.
4. **Analysis agent:** runs code on data, checks statistics, and updates beliefs.
5. **Writing and review agents:** draft papers and critique them before humans see them.

What separates useful systems from paper generators is **verification**: grounding every claim in a source, running real code or experiments, and measuring how often predictions turn out right.

State of the art in 2026

- **Google AI co-scientist:** built on Gemini, with validations reported for acute myeloid leukaemia drug candidates and liver-fibrosis targets [1].
- **FutureHouse Robin:** a multi-agent system whose best-known result is proposing ripasudil, a glaucoma drug, as a candidate for dry age-related macular degeneration; wet-lab work was done by humans [1].
- **Sakana AI Scientist:** produced a paper that scored above the average human acceptance threshold at a workshop, a lower bar than a main conference track [1].
- **Criticism:** commentators argue much AI-generated research is of low quality and novelty, and an early study found chatbots generated few insightful hypotheses in vaccinology [1].

What DARPA is funding

- **expMath (Exponentiating Mathematics):** aims to radically accelerate pure mathematics with an AI co-author that proposes and proves useful abstractions. It has teams on AI methods and on evaluation against professional-level problems; it kicked off in 2025 [2][3].
- **GENIUS (UMD + MIT, 2026):** a US$2.6M expMath effort building a neuro-symbolic reasoning engine that embeds mathematical strategies into language models [4].
- **MATHBAC (2026):** targets next-generation agentic AI platforms that systematically improve themselves to transform scientific discovery for defense [5].
- **DICE (2026):** studies how up to 100,000 AI agents can coordinate without a central orchestrator, the "society of agents" behind large-scale discovery [7].
- Earlier, **ASKE** and **ASKEM** worked on automatically extracting and modeling scientific knowledge, and the **AI Cyber Challenge** showed AI systems finding and fixing software vulnerabilities.

How HAYAT approaches autonomous research

HAYAT, my open-source, local-first assistant, is a practical, small-scale version of these ideas running on one PC [8]:

- **Research with receipts:** background web research cites its sources and checks freshness and conflicts; web content is always treated as untrusted data.
- **Learning gaps:** topics HAYAT couldn't answer can be researched in the background and remembered (with a switch to turn it off).
- **Predictions with track records:** HAYAT keeps measured track records for its predictions instead of just sounding confident.
- **Proof of action:** it only says something was done when a tool result shows it, and outward-facing actions need approval.

- **Verified engineering:** 597 modules and 7,865 automated tests keep the agent loop, memory and tools honest.

The design rule I took from building it: an autonomous system earns trust by showing its work, not by sounding smart.

How it could change life

- **Health:** faster drug repurposing, where an AI finds an approved drug that might treat another disease, can bring treatments to patients years sooner.
- **Materials and energy:** AI-guided discovery of battery chemistries, catalysts and alloys.
- **Mathematics and software:** AI co-authors that suggest lemmas, find counterexamples and optimize algorithms (as AlphaEvolve did).
- **Access:** small teams, startups and universities with limited budgets can run research programs that once needed a large lab.

Future scenarios

2027	2030	2035	2040+
AI co-scientists are standard in pharma and materials R&D, with humans choosing what to test.	Closed-loop self-driving labs in narrow domains: AI proposes, robots test, AI updates.	AI co-authorship is routine in math and science; verification capacity is the bottleneck.	Discovery at "agent-society" scale, with strong provenance standards for machine-generated science.

- **Optimistic:** discovery rates rise sharply in medicine and materials while review systems keep quality high.
- **Base case:** AI dramatically speeds the routine parts of research; breakthroughs still need human insight.
- **Pessimistic:** floods of low-quality AI papers overwhelm peer review and erode trust.

Effects and risks

- **Quality and trust:** cheap papers can drown good ones; journals need provenance rules and AI-use disclosure.
- **Reproducibility:** every AI-generated claim should come with code, data and logs.
- **Dual use:** automated discovery in biology and chemistry needs safeguards and access controls.
- **Credit:** who gets authorship when an AI proposed the idea?

What to watch

- Independent replication of AI-proposed findings in wet labs.
- expMath results on professional-level math problems.
- Journal policies on AI-generated research and reviewer workload.
- Open benchmarks for hypothesis quality, not just paper fluency.

My take

The next decade of AI for science will be won on verification. The systems I trust most will cite every claim, run real experiments, log everything, and say "I don't know" when that's the truth. It's the same

standard I hold HAYAT to, and the same standard an SRE holds a production system to: if you can't observe it and prove it, it didn't happen.

Key players & programs

Google DeepMind	AI co-scientist (Nature, 2026); AlphaEvolve (2025); AlphaFold [1][6]
FutureHouse / Edison Scientific	Robin multi-agent discovery system (Nature, 2026) [1]
Sakana AI	The AI Scientist, end-to-end automated papers (Nature, 2026) [1]
DARPA	expMath, MATHBAC, DICE, and earlier ASKE/ASKEM [2][5][7]
HAYAT (open source)	Local-first assistant with cited research, background learning and verified actions [8]

Frequently asked questions

What is an AI scientist?
A system, usually several AI agents working together, that reads literature, proposes hypotheses, designs experiments or analyses, and writes up results, with humans supervising key steps.

Has AI made real scientific discoveries?
Yes, with caveats. AlphaFold transformed protein-structure prediction, AlphaEvolve improved a matrix-multiplication algorithm, and 2026 AI-scientist papers reported drug-repurposing candidates, though humans did the lab validation.

What is DARPA expMath?
Exponentiating Mathematics, a DARPA program launched in 2025 to build an AI co-author that proposes and proves useful mathematical abstractions.

How does HAYAT relate to autonomous research?
HAYAT is Mohsin Khan Shaik's open-source local AI assistant. It runs cited background research, learns topics it couldn't answer, tracks how accurate its predictions are, and only reports actions that tool results prove.

Glossary

Self-driving lab: A laboratory where AI plans experiments and robots run them in a closed loop.

Provenance: A record of where a claim, dataset or result came from.

Neuro-symbolic AI: Systems that combine neural networks with formal logic or symbolic reasoning.

Drug repurposing: Finding new uses for drugs that are already approved for other conditions.

References

1. Singularity Hub — What the rise of AI scientists may mean for human research (Feb 2026)
   https://singularityhub.com/2026/02/20/what-the-rise-of-ai-scientists-may-mean-for-human-research/
2. DARPA — expMath: Exponentiating Mathematics
   https://www.darpa.mil/research/programs/expmath-exponential-mathematics
3. DARPA — Accelerating mathematical discovery with AI (2025)
   https://www.darpa.mil/news/2025/math-ai-tomorrows-breakthroughs
4. University of Maryland — UMD to lead DARPA-funded effort to accelerate mathematical discovery with AI
   https://today.umd.edu/umd-to-lead-darpa-funded-effort-to-accelerate-mathematical-discovery-with-ai

5. DARPA — MATHBAC program
   https://www.darpa.mil/research/programs/mathbac
6. Google DeepMind — AlphaEvolve: a Gemini-powered coding agent for designing advanced algorithms (May 2025)
   https://deepmind.google/discover/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/
7. DARPA — DICE: Decentralized AI through Controlled Emergence
   https://www.darpa.mil/research/programs/decentralized-artificial-intelligence-through-controlled-emergence
8. HAYAT on GitHub (README, architecture and test counts)
   https://github.com/mohsinkhanshaik/HAYAT
9. Scientific American — New Google AI tackles complex math and science
   https://www.scientificamerican.com/article/new-google-ai-chatbot-tackles-complex-math-and-science/

03 · COMPUTING

Photonic computers: doing AI math with light

A deep dive into optical computing for AI: how light performs matrix math, the 2025 Nature milestones from Lightmatter and Lightelligence, why interconnect is changing first, and what it means for the energy cost of AI.

9 Apr 2025	16,000+	ResNet · BERT	2026
two photonic AI chips in Nature	photonic components on one PACE chip	real models run on photonics	photonic interconnect for AI clusters

Key takeaways

- Photonic computing uses light to perform the matrix multiplications at the heart of neural networks, potentially with far less energy and latency.
- In April 2025, Nature published processors from Lightmatter and Lightelligence that ran real AI workloads, but both are hybrids: electronics still handle much of the work.
- The near-term revolution is optical interconnect: moving data between AI chips with light, which is where power and bandwidth limits bite hardest.
- Precision, light-to-electronics conversion and manufacturing are the main obstacles to all-optical AI.
- For infrastructure teams, photonics is a power and networking story first.

Background

Optical computing has been studied since the 1960s, but electronics kept winning because transistors kept shrinking. Two things changed: AI workloads became dominated by one operation (matrix multiplication), and data centers hit limits on power and on how fast chips can communicate. Light is excellent at both: photons travel without resistive heating and different colors can carry separate data streams in the same waveguide.

A key paper was Shen et al. (2017, *Nature Photonics*), which ran a simple neural network on a mesh of Mach-Zehnder interferometers, work whose authors went on to found Lightmatter. DARPA's PIPES program (Photonics in the Package for Extreme Scalability) later pushed optical I/O directly into chip packages.

How photonic AI works

- **Encoding:** numbers are encoded in the amplitude or phase of light.
- **Computing:** a mesh of tiny interferometers splits and recombines light so the output intensities equal a matrix-vector product, performed as the light passes through, in effect "at the speed of light".
- **Reading out:** photodetectors convert the result back into electrical signals for nonlinear functions, memory and control.
- **Interconnect:** separately, silicon photonics transceivers and co-packaged optics send data between chips over fiber, with lower energy per bit than copper over distance.

State of the art in 2026

- **Lightmatter (Nature, Apr 2025):** a photonic processor that ran state-of-the-art neural networks including ResNet and BERT, plus segmentation and reinforcement-learning tasks [1][2].
- **Lightelligence PACE (Nature, Apr 2025):** a hybrid accelerator integrating more than 16,000 photonic components on one silicon chip, focused on ultra-low latency [2].
- **Hybrid reality:** in both, light performs the core math while electronics do everything else [2].
- **Interconnect (2026):** Lightmatter's *Guide* light engine uses a "very large scale photonics" laser architecture to connect thousands to millions of accelerators [3]; switch vendors are shipping co-packaged optics for AI networking.

Why interconnect comes first

Large AI training runs spread one model across thousands of GPUs that constantly exchange data. Much of the cluster's power and cost goes to moving bits, not computing on them. Optical links reduce energy per bit, reach farther and carry more bandwidth per fiber. That is a problem infrastructure teams face today, which is why optical I/O and co-packaged optics are being adopted before optical compute.

How it could change life

- **Cheaper AI:** lower energy per inference brings down the cost of AI services.
- **Greener data centers:** less electricity per answer eases pressure on grids and water for cooling.
- **Real-time AI:** ultra-low-latency inference for autonomous vehicles, finance and robotics.
- **AI in constrained places:** satellites, telecom base stations and edge devices with limited power.

Future scenarios

2027	2030	2035	2040+
Co-packaged optics common in new AI clusters; optical I/O chiplets in accelerators.	Hybrid photonic accelerators in production for specific inference workloads.	Optical compute is a standard part of the AI stack, alongside GPUs.	Mostly-optical data movement inside and between racks; analog optical math for selected models.

- **Optimistic:** photonic accelerators deliver order-of-magnitude energy gains for inference.
- **Base case:** photonics dominates interconnect while compute stays mostly electronic.
- **Pessimistic:** conversion overheads and precision limits keep optical compute niche.

Effects and risks

- **Precision:** analog optical math is noisy; models must be trained to tolerate it.
- **Conversion overhead:** every trip between light and electronics costs energy and latency.
- **Manufacturing:** photonic fabrication, lasers and packaging are less mature than CMOS.
- **Supply chains:** a new hardware layer brings new dependencies.

What to watch

- Energy-per-token measurements on real models, not just peak TOPS.
- Co-packaged optics in GPU-to-GPU fabrics.

- Training (not just inference) on photonic hardware.
- Standards for optical chiplet interfaces.

Connection to my work

In AI infrastructure, the network and the power budget are often the real limits on cluster size. Photonic interconnect changes capacity planning, failure modes (lasers and optics fail differently from copper) and observability (optical power levels become key metrics). I'd plan for it as an infrastructure migration: pilot racks, telemetry first, then scale.

My take

Bet on photonic interconnect now and treat all-optical AI as a promising long game. The companies that solve packaging and reliability will matter more than those with the flashiest demos.

Key players & programs

Lightmatter	Photonic processor (Nature, 2025); Passage and Guide photonic interconnect [1][3]
Lightelligence	PACE hybrid photonic accelerator, 16,000+ components (Nature, 2025) [2]
DARPA PIPES	Photonics in the package for extreme scalability (optical I/O)
Network and switch vendors	Co-packaged optics for AI cluster networking

Frequently asked questions

What is a photonic computer?
A computer that uses light instead of electrical current to move or process information. For AI, photonic chips perform matrix multiplications as light passes through interferometer meshes.

Are photonic AI chips real today?
Yes as research prototypes and early products. In April 2025 Nature published photonic processors from Lightmatter and Lightelligence that ran real AI models, though both still rely on electronics for parts of the work.

Will photonics replace GPUs?
Not soon. The near-term change is optical interconnect between chips; photonic compute is likely to complement GPUs for specific workloads.

Why does photonics save energy?
Light moves without resistive heating and can carry many data streams at once, reducing energy per operation and per bit moved.

Glossary

Mach-Zehnder interferometer (MZI): An optical component that splits and recombines light to control its output; meshes of MZIs perform matrix math.

Co-packaged optics (CPO): Optical transceivers placed in the same package as a switch or processor chip.

Silicon photonics: Building optical components with standard silicon chip manufacturing.

Inference: Running a trained AI model to produce outputs.

References

1. Science News — New computer chips do math with light (Apr 2025)
   https://www.sciencenews.org/article/computer-chips-math-light-photonic
2. Physics World — Photonic computer chips perform as well as purely electronic counterparts
   https://physicsworld.com/a/photonic-computer-chips-perform-as-well-as-purely-electronic-counterparts-say-researchers/
3. insideHPC — Lightmatter introduces Guide light engine for AI (Jan 2026)
   https://insidehpc.com/2026/01/lightmatter-introduces-guide-light-engine-for-ai-with-vlsp-technology/
4. Lightmatter — A new kind of computer
   https://lightmatter.co/blog/a-new-kind-of-computer/

Online: https://mohsinkhanshaik.com/research/photonic-computing · Facts are cited; outlooks and "My take" are the author's opinion.

04 · QUANTUM

Quantum cybersecurity: the migration that can’t wait

A deep dive into the quantum threat to encryption: how Shor's algorithm breaks RSA, the NIST post-quantum standards, how fast the threat estimates are moving, where QKD fits, and a practical migration playbook.

Aug 2024	60%+	<1M	2035
NIST FIPS 203 / 204 / 205	of human web traffic using hybrid PQ (early 2026)	noisy qubits to break RSA-2048 (2025 estimate)	RSA-2048 disallowed (NIST draft)

Key takeaways

- A large, error-corrected quantum computer could break RSA and elliptic-curve cryptography, which protect most of today's internet.
- Post-quantum cryptography (PQC) is standardized: NIST published FIPS 203, 204 and 205 in August 2024 and picked HQC as a backup in 2025.
- Estimates of the quantum resources needed keep falling: under one million noisy qubits for RSA-2048 (Gidney, 2025), with newer preprints claiming fewer.
- Adoption is real on the client side, over 60% of human web traffic to Cloudflare used hybrid post-quantum key exchange by early 2026, but servers and internal systems lag badly.
- 'Harvest now, decrypt later' means long-lived secrets are already at risk; the migration is an engineering program, not a one-off patch.

Background: why quantum computers threaten encryption

Public-key cryptography relies on math problems that are easy one way and hard to reverse: factoring large numbers (RSA) and the discrete logarithm on elliptic curves (ECDH, ECDSA). In 1994 Peter Shor showed a quantum computer could solve both efficiently. Symmetric encryption (AES) and hashing are affected far less: Grover's algorithm only gives a square-root speed-up, which longer keys offset.

The danger isn't hypothetical just because the machine doesn't exist yet. Adversaries can record encrypted traffic today and decrypt it once a capable quantum computer arrives, known as **"harvest now, decrypt later."** Anything that must stay secret for 10–20 years is already exposed.

The standards

- **FIPS 203 (ML-KEM):** a lattice-based key-encapsulation mechanism (formerly CRYSTALS-Kyber) for establishing shared keys.
- **FIPS 204 (ML-DSA):** a lattice-based digital signature (formerly CRYSTALS-Dilithium).
- **FIPS 205 (SLH-DSA):** a hash-based signature (SPHINCS+) as a conservative backup.
- **HQC:** selected in March 2025 as a code-based backup KEM; FIPS 206 (FN-DSA, Falcon) is in development [1][2].
- **Transition:** NIST's draft IR 8547 deprecates 112-bit algorithms such as RSA-2048 after 2030 and disallows them after 2035 [2].

How fast is the threat moving?

- **2019:** about 20 million noisy qubits were estimated to factor RSA-2048 in hours.
- **May 2025:** Gidney estimated fewer than one million noisy qubits in under a week [3].
- **2026:** a preprint claims under 100,000 physical qubits using quantum LDPC codes under specific hardware assumptions, and a Google paper put elliptic-curve secp256k1 (used by Bitcoin and Ethereum) at around 500,000 physical qubits [2].
- These are estimates under assumptions, not demonstrations; today's largest machines have hundreds to low thousands of physical qubits.

Adoption in 2026

- **Clients:** Cloudflare reported client support for hybrid post-quantum key exchange rising from under 3% in early 2024 to over 60% of human traffic by February 2026 [4].
- **Servers:** only about 10–15% of origin servers Cloudflare connects to support it [4]. OpenSSL 3.5, GnuTLS 3.8.9 and Go 1.24 enabled hybrid PQ key exchange by default in 2025 [4].
- **Messaging:** Signal (PQXDH) and Apple iMessage (PQ3) added post-quantum protection for messages.
- **Big tech:** Google targets 2029 for full post-quantum security, including authentication [2].

Where QKD fits

Quantum key distribution uses quantum physics to detect eavesdropping while sharing keys. China has operated satellite QKD since the Micius satellite (2016), and ESA's Eagle-1 is due around 2026–27 [6]. But the NSA still does not recommend QKD for sensitive US systems: it requires special hardware and dedicated links, doesn't authenticate the source by itself, and is hard to validate [5]. For almost every organization, PQC is the practical path; QKD is a niche complement.

A practical migration playbook

1. **Inventory:** find every place cryptography lives: TLS, SSH, VPNs, code signing, tokens and JWTs, KMS/HSMs, databases, IoT firmware and third-party SaaS.
2. **Prioritize by secrecy lifetime:** data that must stay secret beyond 2030 goes first.
3. **Build crypto-agility:** make algorithms configuration, not code, so they can be swapped without redeploying everything.
4. **Roll out hybrid first:** combine classical and PQ (for example X25519 + ML-KEM) so security never drops below today's level.
5. **Measure:** watch handshake latency, packet sizes, error rates and middlebox breakage; roll out behind feature flags with canaries.
6. **Signatures and PKI:** plan certificate chains with PQ signatures; they are much larger and affect bandwidth and storage.
7. **Vendors:** add PQ readiness to procurement and security reviews.

How it could change life

Done well, users notice nothing: browsers, phones and apps quietly switch to quantum-safe cryptography. Done badly, we could see a sudden loss of trust in banking, health records, elections and

cryptocurrencies when a capable quantum computer appears. The migration protects privacy for decades of data that is being collected right now.

Future scenarios

2027	2030	2033–2035	Q-day
Hybrid PQ key exchange is default across browsers, clouds and major VPNs.	RSA-2048 deprecated; regulated sectors complete inventories and most migrations.	Government-wide transition deadlines; classical public-key crypto disallowed for most federal use.	Unknown. Expert estimates range from around 2029 (aggressive) to well after 2035.

- **Optimistic:** migration completes before a cryptographically relevant quantum computer exists.
- **Base case:** the internet edge is safe, but legacy and embedded systems remain exposed for years.
- **Pessimistic:** an early Q-day exposes harvested data and unmigrated systems, including blockchains.

Effects and risks

- **Migration risk:** rushed rollouts break things; most inventories are incomplete.
- **Performance:** larger keys and signatures affect handshakes, constrained IoT and certificate sizes.
- **Blockchains:** elliptic-curve signatures need coordinated, community-wide upgrades.
- **Inequality:** smaller organizations and countries may lag, becoming easy targets.

What to watch

- Error-corrected logical-qubit counts on real hardware.
- Server-side (origin) PQ adoption rates.
- Final NIST IR 8547 and sector-specific deadlines.
- PQ signatures in the Web PKI and certificate transparency.

My take

This is the largest infrastructure migration most organizations will run this decade, and it should be run like one: inventory, SLOs, canaries, dashboards and rollback plans. The math is settled; the operations are not. Start with long-lived secrets and build crypto-agility so the next algorithm change is a config change.

Key players & programs

NIST	FIPS 203/204/205 (2024), HQC (2025), transition plan IR 8547 [1][2]
NSA	CNSA 2.0 guidance; does not recommend QKD for sensitive US systems [5]
Cloudflare	Measures and deploys hybrid PQ TLS at internet scale [4]
Google	2029 target for full post-quantum security; resource estimates [2][3]
China / ESA	Satellite QKD: Micius (2016); Eagle-1 planned [6]

Frequently asked questions

What is post-quantum cryptography?
Encryption and signature algorithms designed to resist both classical and quantum computers, such as NIST's ML-KEM, ML-DSA and SLH-DSA standards from 2024.

When will quantum computers break RSA?
No one knows. Resource estimates have fallen to under one million noisy qubits for RSA-2048, but no machine is close today; estimates for 'Q-day' range from around 2029 to well after 2035.

What is 'harvest now, decrypt later'?
Recording encrypted data today in order to decrypt it once a capable quantum computer exists. It makes long-lived secrets vulnerable now.

Is quantum key distribution the answer?
For most organizations, no. The NSA does not recommend QKD for sensitive US systems because of hardware, authentication and validation limits. Post-quantum cryptography is the practical path.

Glossary

Shor's algorithm: A quantum algorithm that factors integers and computes discrete logarithms efficiently, breaking RSA and ECC.

ML-KEM: Module-Lattice-Based Key-Encapsulation Mechanism, standardized as FIPS 203.

Hybrid key exchange: Combining a classical and a post-quantum algorithm so both must be broken.

Crypto-agility: The ability to change cryptographic algorithms quickly, by configuration.

References

1. NIST — First 3 finalized post-quantum encryption standards (Aug 2024)
   https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards
2. The Quantum Insider — Why 2026 matters for quantum security
   https://thequantuminsider.com/2026/04/28/why-2026-matters-quantum-security/
3. Gidney — How to factor 2048-bit RSA integers with less than a million noisy qubits (arXiv, 2025)
   https://arxiv.org/abs/2505.15917
4. Cloudflare — Origin PQ readiness, key transparency and ASPA on Radar (Feb 2026)
   https://blog.cloudflare.com/radar-origin-pq-key-transparency-aspa/
5. NSA — Quantum Key Distribution (QKD) and Quantum Cryptography
   https://www.nsa.gov/Cybersecurity/Quantum-Key-Distribution-QKD-and-Quantum-Cryptography-QC/
6. GMF — Quantum entangles the heavens (satellite QKD)
   https://www.gmfus.org/news/quantum-entangles-heavens
7. Cloudflare Radar — Post-quantum encryption worldwide (live)
   https://radar.cloudflare.com/post-quantum

05 · DARPA

Atomically precise: DARPA APM and the road from atoms to products

A deep dive into atomically precise manufacturing: from Feynman's 1959 vision and DARPA's 2015 Atoms to Product program to the 2026 Atomically Precise Molecules program and its programmable precision polymers.

1959	2015	36 months	Apr 2027
Feynman: 'Plenty of Room at the Bottom'	Atoms to Product: 10 performers	APM program, two phases	assumed APM start

Key takeaways

- Atomically precise manufacturing means building materials and devices where the position of every atom or monomer is controlled.
- DARPA's Atoms to Product (A2P, 2015) tried to carry nanoscale properties up to millimeter-scale products.
- DARPA's Atomically Precise Molecules (APM, 2026) targets programmable precision polymers, controlling sequence, orientation and composition, with nonlinear optical materials as a first target.
- Biology already does this: proteins and DNA are sequence-defined polymers; APM aims to bring that control to synthetic chemistry at scale.
- Scale-up and cost are the main challenges; AI-driven design plus automated chemistry is the likely path.

Background

In 1959 Richard Feynman's lecture "There's Plenty of Room at the Bottom" imagined arranging atoms one by one. Decades later, scanning tunneling microscopes made it literal: IBM researchers spelled their company's name with 35 xenon atoms in 1989. In 1986 Eric Drexler's *Engines of Creation* popularized molecular manufacturing. Since then, chemistry has moved toward precision from the other direction: DNA origami, sequence-defined polymers and engineered proteins.

Atoms to Product (2015)

DARPA's A2P program aimed to develop ways to assemble atomic-scale pieces into larger products while keeping their nanoscale properties, closing the "assembly gap" in two steps: from atoms to microns, then microns to millimeters [1][2]. DARPA selected 10 performers in December 2015, including work on top-down, atomically precise fabrication [2]. The lesson from A2P was that precision is possible at small scales but expensive and slow to scale up.

APM: Atomically Precise Molecules (2026)

- **Goal:** synthesis approaches that deliver **programmable precision polymers**, controlling the sequence, orientation and composition of molecules [3].

- **First target:** nonlinear optical materials, materials whose optical behavior changes with light intensity and that matter for lasers, sensing and optical communication [4].
- **Structure:** a two-phase, 36-month effort: Phase 1 develops new chemistries, Phase 2 applies them to materials and devices, with demonstration milestones through the program [3][4].
- **Logistics:** solicitation DARPA-PS-26-141 published 25 September 2026, proposers day 30 September 2026, abstracts due 27 October 2026; program manager John M. Hoffman in DARPA's Multi X Office [3] [4].
- **Open to biology:** DARPA's Q&A indicates biological or biochemical approaches are welcome if they meet the goals and metrics [4].

How precision polymers work

Most plastics are made by chain reactions that join monomers randomly, giving chains of varying length and order. **Sequence-defined polymers** place each monomer in a chosen order, the way a ribosome builds a protein. Approaches include iterative step-by-step synthesis on solid supports, templated synthesis (using DNA or another template), and engineered enzymes that accept non-natural building blocks. Precise sequence means precise properties: how the chain folds, which molecules it binds, how it conducts light or charge.

How it could change life

- **Water and environment:** membranes that filter specific pollutants or extract lithium from brine.
- **Energy:** catalysts that waste less energy, better battery electrolytes and separators.
- **Medicine:** drug-delivery materials that release exactly where and when needed.
- **Optics and communication:** nonlinear optical materials for faster, lower-power photonic devices, linking APM to photonic computing.

Future scenarios

2028
New chemistries that place monomers in defined sequences at useful quantities.

2030
First devices using precision polymers: optical materials, sensors, membranes.

2035
AI-designed sequences made by automated chemistry, "compile a material".

2040+
Programmable materials manufactured at industrial scale for selected high-value uses.

- **Optimistic:** precision polymers become a design platform like software, with libraries of reusable sequences.
- **Base case:** high-value niches (optics, medical, separation) adopt them first.
- **Pessimistic:** cost and scale keep precision chemistry in the lab.

Effects and risks

- **Scale:** grams in the lab versus tons in industry.
- **Cost:** step-by-step synthesis is slow and expensive.
- **Safety:** new materials need toxicology and end-of-life planning.
- **Dual use:** precise chemistry can enable both protective and harmful materials; oversight matters.

What to watch

- APM awards and first-phase chemistry results.
- Throughput (monomers per hour) and yield of precise synthesis.
- AI models that predict polymer properties from sequence.
- Integration of APM optical materials into photonic devices.

My take

APM is a sign that materials science is becoming programmable. The pattern I expect mirrors software delivery: a specification (the sequence), an automated build (synthesis), and automated tests (characterization), with AI in the loop. The bottleneck will be throughput and observability of the "build", which is exactly the kind of pipeline problem engineers know how to attack.

Key players & programs

DARPA APM (2026)	Programmable precision polymers; first target nonlinear optical materials [3] [4]
DARPA A2P (2015)	Assembling nanoscale pieces into millimeter-scale products; 10 performers [1][2]
Zyvex	A2P performer working on top-down atomically precise fabrication [2]

Frequently asked questions

What is DARPA APM?
Atomically Precise Molecules, a 2026 DARPA program to develop synthesis methods for programmable precision polymers, controlling sequence, orientation and composition, with nonlinear optical materials as a first target.

What was DARPA Atoms to Product?
A 2015 DARPA program to assemble atomic-scale pieces into millimeter-scale products while keeping nanoscale properties. It selected 10 performers in December 2015.

What are precision polymers?
Polymers where each monomer is placed in a defined order, like a protein, rather than joined randomly as in most plastics.

Why does atomic precision matter?
Precise structure gives precise properties, enabling better membranes, catalysts, optical materials and drug-delivery systems.

Glossary

Monomer: A small molecule that links with others to form a polymer.

Sequence-defined polymer: A polymer whose monomers are arranged in a specific, controlled order.

Nonlinear optical material: A material whose optical response depends on light intensity, used in lasers and optical signal processing.

Scanning tunneling microscope: An instrument that images and manipulates individual atoms on surfaces.

References

1. DARPA — Atoms to Product program page
   https://www.darpa.mil/research/programs/atoms-to-product
2. DARPA — Program seeks ability to assemble atom-sized pieces into practical products (2015)
   https://www.darpa.mil/news/2015/assemble-atom-sized-pieces
3. DARPA — APM: Atomically Precise Molecules
   https://www.darpa.mil/research/programs/apm
4. DARPA — APM proposers day presentation (Oct 2026)
   https://www.darpa.mil/sites/default/files/attachment/2026-10/darpa-program-apm-proposers-day-presentation.pdf
5. DARPA — APM Q&A (Oct 2026)
   https://www.darpa.mil/sites/default/files/attachment/2026-10/program-apm-q-a.pdf

06 · DARPA

DARPA DICE: 100,000 AI agents and no boss

A deep dive into decentralized AI: why today's multi-agent systems still need an orchestrator, what DARPA's DICE program will measure, how 'controlled emergence' works, and what it means for resilient infrastructure.

100,000	1 million	4	Apr 2026
agents in final simulations	messages exchanged	technical fields: AI, control, formal methods, game theory	program announced

Key takeaways

- DICE (Decentralized Artificial Intelligence through Controlled Emergence) seeks theory, algorithms and architectures for coordinating very large AI agent collectives without a central controller.
- It is simulation-only: DARPA states it does not cover deploying autonomous systems in the real world.
- The final phase targets simulated collectives of up to 100,000 agents exchanging up to 1 million messages.
- Success is measured with metrics such as number of agents, interactions, failures and compromises, and mission success rate, using a common simulation environment.
- 'Controlled' means statistical guarantees of staying aligned with the mission, not strict determinism.

Background: why decentralize?

Most multi-agent AI systems today, from LLM agent frameworks to robot fleets, rely on a central orchestrator that assigns tasks and merges results. That works at small scale but creates a bottleneck and a single point of failure, and it breaks when communication is cut. Nature does it differently: ant colonies, flocks and immune systems coordinate through local interactions. DARPA explored this before with programs such as OFFSET (swarms of drones in urban environments); DICE asks the question for AI agents at a much larger scale.

What DICE asks for

- **Goal:** the underlying theory, algorithms and architectures for decentralized coordination among large numbers of AI agents [1].
- **Disciplines:** artificial intelligence, control theory, formal methods and game theory [3].
- **Scale:** final-phase simulations of up to 100,000 agents exchanging up to 1 million messages [4].
- **Evaluation:** a test-and-evaluation team runs a common simulation environment and scores metrics including number of agents, interactions, failures and compromises, and mission success rate [3].
- **Scope:** simulation-centered; no real-world deployment of autonomous systems [1].
- **Logistics:** announced in April 2026; BAA HR001126S0010 with proposals due 25 August 2026; program manager Susmit Jha [2][3].

What 'controlled emergence' means

Emergence is when simple local rules produce complex group behavior. It is powerful and dangerous: the same mechanisms that create flocking can create stampedes. DARPA defines control in statistical terms: it expects stochastic convergence with statistical guarantees, keeping the collective aligned with its mission and agents in coherent roles within bounds, rather than strict determinism [3]. In engineering terms: provable envelopes, not scripts.

The distributed-systems view

Everything distributed-systems engineers know applies at agent scale:

- **Consensus and gossip:** how agents agree on state without a leader.
- **Backpressure and rate limits:** preventing message storms in a million-message system.
- **Failure domains:** limiting how far a bad agent or partition can spread its effects.
- **Byzantine fault tolerance:** staying correct even when some agents are compromised and lie.
- **Observability:** tracing decisions across thousands of agents to explain outcomes.

How it could change life

- **Disaster response:** logistics agents that keep coordinating when networks are damaged.
- **Energy grids:** millions of devices (batteries, EVs, heat pumps) balancing supply and demand locally.
- **Software operations:** fleets of agents that detect, diagnose and fix incidents without a single control plane.
- **Science:** large populations of research agents exploring hypotheses in parallel.

Future scenarios

2027	2029	2032	2035+
Theory and benchmarks for controlled emergence; first scaled simulations.	Open protocols for agent-to-agent coordination with statistical safety bounds.	Decentralized agent networks in logistics and infrastructure operations.	Agent "societies" managing complex systems under formal guarantees and human oversight.

- **Optimistic:** provable decentralized coordination becomes a standard layer for AI systems.
- **Base case:** hybrid designs: decentralized execution with central policy and audit.
- **Pessimistic:** emergent failures make large agent collectives too risky outside simulation.

Effects and risks

- **Emergent failure:** cascading behavior no single agent intended, like flash crashes.
- **Accountability:** who is responsible when no one is in charge?
- **Security:** compromised agents can poison many others.
- **Military use:** DARPA's framing is defense; clear human-control policies are essential.

What to watch

- DICE awards and published theory on controlled emergence.
- Benchmarks: mission success rate versus compromised-agent fraction.

- Open-source agent-coordination protocols adopting formal guarantees.

Connection to my work

I built an eBPF self-healing agent with exactly these guardrails: thresholds over windows, cooldowns, hourly budgets, one in-flight action per rule, and a circuit breaker that stops acting and pages a human. Multiply that agent by 100,000 and you have DICE's problem. Budgets, circuit breakers and blast-radius limits will matter as much as clever algorithms.

My take

DICE is the most SRE-like program in DARPA's portfolio. The teams that win will treat agent collectives like production distributed systems: measure everything, assume failure, design for partitions and adversaries, and keep humans able to stop the system.

Key players & programs

DARPA DICE	Decentralized coordination of up to 100,000 simulated agents [1][4]
DARPA OFFSET (2017–2021)	Earlier swarm program for drones and ground robots in urban settings
Multi-agent AI frameworks	Today's LLM agent systems, mostly orchestrator-based

Frequently asked questions

What is DARPA DICE?
Decentralized Artificial Intelligence through Controlled Emergence, a 2026 DARPA program developing theory, algorithms and architectures for coordinating very large AI agent collectives without a central controller.

How many agents will DICE simulate?
Its final phase targets simulated collectives of up to 100,000 agents exchanging up to 1 million messages.

Will DICE deploy autonomous systems?
No. DARPA states the program is simulation-centered and does not cover developing or launching autonomous systems in the real world.

What does 'controlled emergence' mean?
Letting useful group behavior emerge from local agent interactions while keeping the collective aligned with its mission within statistical bounds.

Glossary

Orchestrator: A central component that assigns tasks to agents and combines their results.

Emergence: Complex group behavior arising from simple local interactions.

Byzantine fault: A failure where a component behaves arbitrarily or maliciously, including sending false information.

Gossip protocol: A way for nodes to share information by repeatedly exchanging it with random peers.

References

1. DARPA — DICE program page
https://www.darpa.mil/research/programs/decentralized-artificial-intelligence-through-controlled-emergence

2. Intelligence Community News — DARPA discusses DICE program opportunity
   https://intelligencecommunitynews.com/darpa-discusses-dice-program-opportunity/
3. DARPA — DICE Q&A (Aug 2026)
   https://www.darpa.mil/sites/default/files/attachment/2026-08/darpa-program-dice-q-a.pdf
4. The Debrief — DARPA wants to build an army of 100,000 AI agents
   https://thedebrief.org/darpa-wants-to-build-an-army-of-100000-ai-agents-that-can-think-and-act-on-their-own-dice/
5. DARPA — DICE proposers day (2026)
   https://www.darpa.mil/sites/default/files/attachment/2026-06/dice-proposers-day-2026.pdf

07 · DARPA

How DARPA research changes the world — and what's next

A deep dive into DARPA: how a small agency founded after Sputnik helped create the internet and GPS, how its program-manager model works, what's in the 2026 portfolio, and what engineers can learn from it.

1958	US$4.1B	6	4
founded after Sputnik	FY2024 enacted budget	technical offices	companies in QBI Stage C (Oct 2026)

Key takeaways

- DARPA funds high-risk, high-reward research through temporary program managers with clear goals, fixed timelines and independent evaluation.
- Its track record includes ARPANET, early satellite navigation, stealth, autonomous vehicles and early mRNA vaccine work.
- The 2026 portfolio spans quantum benchmarking (QBI), AI that does science (expMath, MATHBAC), decentralized AI (DICE), biocomputing (O-CIRCUIT), non-invasive BCIs (Firefox) and atomically precise chemistry (APM).
- QBI is a model for cutting through hype: independent teams test whether a utility-scale quantum computer can be built by 2033.
- The Heilmeier Catechism, DARPA's checklist for new programs, is a useful tool for any engineering proposal.

Origins

The Soviet launch of Sputnik in October 1957 shocked the United States. In February 1958 the Advanced Research Projects Agency (ARPA) was created so the US would never again be surprised by technology. It became DARPA in 1972. Its mission is to make pivotal investments in breakthrough technologies for national security [1].

How DARPA works

- **Program managers:** technical experts hired for a few years each, who design and run programs; turnover keeps ideas fresh.
- **Programs, not labs:** DARPA owns almost no labs; it funds universities, companies and government labs through competitive solicitations.
- **Clear goals and phases:** programs have measurable targets and go/no-go gates; teams that miss them are cut.
- **Independent test and evaluation:** separate teams measure results, as in DICE and QBI.
- **Budget:** about US$4.1 billion enacted for FY2024, with a FY2025 request of US$4.369 billion, across six technical offices [1].

The Heilmeier Catechism

George Heilmeier, DARPA director from 1975 to 1977, created a set of questions every program must answer. They are worth using for any project:

1. What are you trying to do? Articulate your objectives with no jargon.
2. How is it done today, and what are the limits of current practice?
3. What is new in your approach, and why do you think it will succeed?
4. Who cares? If you succeed, what difference will it make?
5. What are the risks?
6. How much will it cost?
7. How long will it take?
8. What are the mid-term and final "exams" to check for success?

A track record

- **ARPANET (1969):** the packet-switched network that became the internet.
- **Satellite navigation:** early work that fed into GPS.
- **Stealth:** the research behind low-observable aircraft.
- **Grand Challenges (2004–2007):** autonomous vehicle competitions that seeded today's self-driving industry.
- **mRNA:** early investments in rapid vaccine platforms that later mattered during COVID-19.
- **AI Cyber Challenge (2023–2025):** AI systems that find and fix software vulnerabilities.

The 2026 portfolio

- **Quantum Benchmarking Initiative (QBI):** will a utility-scale quantum computer be possible by 2033? DARPA selected 11 companies for Stage B in November 2025; on 7 October 2026 Atom Computing, Diraq, IBM and IonQ advanced to Stage C [2][3].
- **AI for discovery:** expMath (AI co-author for mathematics) and MATHBAC (self-improving agentic discovery) [4].
- **DICE:** decentralized coordination of up to 100,000 AI agents in simulation [5].
- **O-CIRCUIT:** living neurons and organoids as a computing medium [6].
- **Firefox:** a 100-channel non-invasive brain-computer interface [7].
- **APM:** programmable precision polymers [8].

How DARPA research changes life

Most of DARPA's impact arrives years later and indirectly: the internet you use, the navigation in your phone, the voice assistant on your desk, and vaccines developed in months instead of years all trace back partly to DARPA bets. Today's programs aim at the next layer: trustworthy quantum computers, AI that accelerates science, safer brain interfaces and new materials.

Future scenarios

2027–2028
QBI Stage C results; first O-CIRCUIT, DICE and Firefox prototypes.

2030
AI-for-science programs expand from math into physics, chemistry and biology.

2033
QBI's target for judging whether utility-scale quantum computing is achievable.

2035+
Spin-offs from today's programs reach civilian products.

Effects and risks

- **Dual use:** defense funding of AI, neurotech and biotech raises ethical questions and needs transparency.
- **Hype vs. evidence:** programs like QBI show how independent evaluation keeps claims honest.
- **Time lag:** civilian benefits may take a decade or more.

What to watch

- QBI Stage C evaluations and any new Stage C entrants.
- Awards under O-CIRCUIT, DICE, Firefox and APM.
- DARPA's FY2026–27 budget priorities.

My take

DARPA's real invention is a management model: bold goal, hard metrics, short tenure, independent evaluation, and permission to fail. It maps neatly onto how good infrastructure teams run reliability programs. I use the Heilmeier questions myself before starting a big project.

Key players & programs

DARPA offices	Six technical offices including Biological Technologies (BTO) and Information Innovation (I2O) [1]
QBI Stage C (Oct 2026)	Atom Computing, Diraq, IBM, IonQ [2]
Universities and labs	Most DARPA research is performed by funded external teams

Frequently asked questions

What is DARPA?
The US Defense Advanced Research Projects Agency, founded in 1958 after Sputnik to fund breakthrough technologies for national security.

What has DARPA invented?
DARPA funded ARPANET (the internet's predecessor), early satellite navigation, stealth research, autonomous vehicle challenges and early mRNA vaccine platforms, among others.

What is the Heilmeier Catechism?
A list of questions created by DARPA director George Heilmeier that every proposed program must answer: goals, current limits, what's new, who cares, risks, cost, time and how success is measured.

What is the DARPA Quantum Benchmarking Initiative?
A program testing whether a utility-scale quantum computer can be built by 2033. In October 2026, Atom Computing, Diraq, IBM and IonQ advanced to Stage C.

Glossary

Program manager (PM): A DARPA technical leader who designs and runs a research program for a few years.

BAA: Broad Agency Announcement, a public solicitation for research proposals.

Proposers day: A public briefing where DARPA explains a new program to potential bidders.

Utility-scale quantum computer: A quantum computer whose computational value exceeds its cost.

References

1. DARPA — About DARPA
   https://www.darpa.mil/about-us/about-darpa
2. DARPA — Quantum Benchmarking Initiative
   https://www.darpa.mil/research/programs/quantum-benchmarking-initiative
3. DARPA — QBI Stage B selection
   https://www.darpa.mil/research/programs/quantum-benchmarking-initiative/stage-b-selection
4. DARPA — expMath
   https://www.darpa.mil/research/programs/expmath-exponential-mathematics
5. DARPA — DICE
   https://www.darpa.mil/research/programs/decentralized-artificial-intelligence-through-controlled-emergence
6. DARPA — O-CIRCUIT
   https://www.darpa.mil/research/programs/o-circuit
7. DARPA — Firefox
   https://www.darpa.mil/research/programs/firefox
8. DARPA — APM
   https://www.darpa.mil/research/programs/apm

Online: https://mohsinkhanshaik.com/research/darpa-research · Facts are cited; outlooks and "My take" are the author's opinion.

08 · BRAIN & NEUROTECH

DARPA Firefox: reading the brain with light, no surgery

A deep dive into non-invasive brain-computer interfaces: why surgery has been needed, how optical interferometry and analog computing could change that, what DARPA's Firefox program targets, and the privacy questions it raises.

100	50 µm	100 µs	36 months
channels, non-invasive	spatial resolution goal	temporal resolution goal	three 12-month phases

Key takeaways

- Firefox is a DARPA Biological Technologies Office program to build a 100-channel non-invasive BCI with 50-micron spatial and 100-microsecond temporal resolution.
- Its two key ideas: dynamic optical scanners that read multiple brain regions at once, and analog computing that processes the data instantly.
- It builds on research detecting neuron firing in human cortex with coherent optical interferometry through a fiber on the scalp.
- Today's best BCIs need brain surgery; a non-invasive device at this resolution would be a major leap.
- Easy-to-wear brain reading makes mental privacy and consent urgent policy questions.

Background: the non-invasive problem

Brain-computer interfaces read neural activity and turn it into commands. Implants like Neuralink's or Utah arrays sit in or on the brain and record individual neurons, but require surgery. Non-invasive methods such as **EEG** (electrodes on the scalp), **fNIRS** (near-infrared light measuring blood oxygen), and **MEG** (magnetic fields) are safe, but the skull blurs signals and blood-flow measures are slow. DARPA's earlier N3 program (Next-Generation Nonsurgical Neurotechnology, 2018) pushed in this direction; Firefox targets far higher resolution.

What Firefox targets

- **Goal:** a 100-channel non-invasive BCI with 50-micron spatial and 100-microsecond temporal resolution [1][2].
- **Innovation 1:** dynamic optical scanners capable of reading multiple brain regions simultaneously [1].
- **Innovation 2:** novel analog computing to process the massive data streams instantly [1].
- **Structure:** 36 months in three 12-month phases with two parallel technical areas; solicitation DARPA-PS-26-137; program manager Dr. Pedro Irazoqui; proposers day in September 2026 [2][3].

How optical brain reading works

When neurons fire, tiny changes in tissue (cell swelling, refractive index, scattering) affect how light travels through it. Coherent optical interferometry compares light that passed through tissue with a reference beam; small changes in phase reveal activity. DARPA notes prior work that extracted the phase

velocity of a small population of neurons firing in human cortex from a non-invasive fiber on the scalp [2]. Scaling that to 100 channels needs fast beam steering (scanners), massive parallel detection, and processing at the sensor, hence analog computing.

Why analog computing

High-resolution optical signals at microsecond timing produce enormous data rates. Digitizing and shipping everything to a computer costs power and latency. Analog front-ends can filter and extract features directly from the signal before digitizing, the same edge-versus-cloud trade-off seen in AI infrastructure.

How it could change life

- **Communication:** people with paralysis or ALS could type or speak through a wearable device without surgery.
- **Diagnosis:** a new window on brain function for epilepsy, stroke recovery and neurodegenerative disease.
- **Rehabilitation:** real-time feedback that helps retrain the brain after injury.
- **Interaction:** hands-free control of computers, vehicles and robots.

Future scenarios

2027	2029	2032	2035+
Phase 1: scanner and analog-processing prototypes in the lab.	Integrated systems approach the 100-channel target in human studies.	Clinical research devices for communication and diagnosis.	Wearable optical BCIs, if size, cost and robustness follow.

- **Optimistic:** near-implant performance without surgery transforms assistive technology.
- **Base case:** meaningful gains over EEG/fNIRS for research and clinical diagnostics.
- **Pessimistic:** scattering through skull and scalp limits resolution far below the target.

Effects and risks

- **Mental privacy:** non-invasive means easy to deploy and potentially easy to misuse; neural data needs strong legal protection and consent.
- **Physics:** light scatters heavily in tissue; hitting 50-micron resolution non-invasively is extremely hard.
- **Equity:** affordable wearables could democratize BCIs, but only if medical access is prioritized.
- **Dual use:** defense interest in human–machine teaming raises ethical questions.

What to watch

- Firefox awards and Phase 1 results.
- Peer-reviewed demonstrations of optical single-neuron-scale signals through the skull.
- Neural-data privacy laws at state and federal level.

My take

Firefox combines two bets I find compelling: optics for sensing and analog compute at the edge. Even partial success would produce valuable spin-offs in medical imaging and sensor processing. The policy

work on neural privacy should start now, before the devices arrive.

Key players & programs

DARPA BTO Firefox	100-channel non-invasive BCI program (2026) [1][2]
DARPA N3 (2018)	Earlier nonsurgical neurotechnology program
Research groups	Optical interferometry, fNIRS, OPM-MEG and functional ultrasound for brain sensing

Frequently asked questions

What is DARPA Firefox?
A DARPA program to build a 100-channel non-invasive brain-computer interface with 50-micron spatial and 100-microsecond temporal resolution, using optical scanners and analog computing.

Is Firefox related to the Firefox web browser?
No. Firefox is the name of a 2026 DARPA neurotechnology program in the Biological Technologies Office.

How can light read the brain?
Neural activity slightly changes how light travels through tissue; coherent optical interferometry detects those tiny changes, even through a fiber placed on the scalp.

Why is non-invasive important?
Today's highest-performance BCIs need brain surgery. A wearable device with similar detail would be safer and far more widely usable.

Glossary

EEG: Electroencephalography: measuring electrical activity with electrodes on the scalp.

fNIRS: Functional near-infrared spectroscopy: measuring blood-oxygen changes with light.

Interferometry: Measuring tiny changes by comparing the phase of two light beams.

Analog computing: Processing continuous signals directly, without converting them to digital numbers first.

References

1. DARPA — Firefox program page
   https://www.darpa.mil/research/programs/firefox
2. DARPA — Firefox FAQ (Oct 2026)
   https://www.darpa.mil/sites/default/files/attachment/2026-10/darpa-program-firefox-faq.pdf
3. DARPA-SN-26-112 Amendment 01 (proposers day)
   https://everglade.com/wp-content/uploads/DARPA-SN-26-112-Amendment-01.pdf

09 · BRAIN & NEUROTECH

Neuralink and the brain-machine interface race

A deep dive into brain-machine interfaces: how implants decode thoughts into actions and speech, Neuralink's progress since its first human implant, the competitors, what peer-reviewed results show, and what remains hype.

Jan 2024	1,024	21+	~97%
first Neuralink human implant	electrodes on 64 threads (N1)	Neuralink participants (early 2026)	speech-BCI word accuracy (NEJM 2024)

Key takeaways

- BMIs record neural activity and decode it into cursor movement, typing or speech; they are already restoring communication for people with paralysis and ALS.
- Neuralink's N1 implant uses 1,024 electrodes on 64 ultra-thin threads placed by a surgical robot; it had 21+ participants by early 2026 and trials in several countries.
- Its vision implant, Blindsight, has FDA Breakthrough Device designation, but no publicly registered human trial had shown vision restoration as of late 2026.
- Competitors are strong: academic speech BCIs reached about 97% word accuracy (NEJM 2024), Precision Neuroscience got FDA 510(k) clearance in 2025, and Synchron and Paradromics continue trials.
- The most important metrics are safety and multi-year reliability, followed by neural-data privacy.

How a BMI works

1. **Sensing:** electrodes record the electrical activity of neurons (spikes and local field potentials).
2. **Signal processing:** filters and spike detection reduce raw data.
3. **Decoding:** machine-learning models map neural patterns to intended movement, letters, words or sounds.
4. **Output:** a cursor moves, text appears, or a synthesized voice speaks; some systems also stimulate the brain to provide feedback.

Designs trade invasiveness for signal quality: penetrating electrodes (Neuralink, Utah arrays) record single neurons; surface films (Precision) sit on the cortex; endovascular stents (Synchron) record from a blood vessel without open-brain surgery.

Neuralink: what's been achieved

- **Device:** the N1 implant with 1,024 electrodes on 64 threads, inserted by the R1 surgical robot, charged wirelessly.
- **First patient (Jan 2024):** Noland Arbaugh, paralyzed from the shoulders down, controlled a computer cursor and played games by thought. Some threads retracted in the weeks after surgery, reducing signal until software changes compensated.
- **Scale:** 21 participants by early 2026, with trials expanding to the UK, Canada and the UAE [1].

- **Speech:** a trial aims to restore speech by decoding speech areas; conversational speed has been described as a target, not yet a published result [1].
- **Blindsight:** FDA Breakthrough Device designation (2024); public trial registries did not show a vision-restoration trial as of late 2026 [2].

The wider field

- **Academic speech BCIs:** a UC Davis team decoded attempted speech for a man with ALS at about 97% word accuracy, published in NEJM in 2024 [3].
- **BrainGate / Blackrock Neurotech:** Utah-array research since the 2000s, with long-term clinical studies.
- **Synchron:** the Stentrode is implanted through a blood vessel, trading some signal fidelity for much lower surgical risk.
- **Precision Neuroscience:** a thin-film cortical interface (Layer 7) with FDA 510(k) clearance in April 2025, used in many patients in clinical settings [4].
- **Paradromics:** a high-bandwidth Connexus implant with early human studies.

How it could change life

- **Independence:** people with paralysis, ALS or stroke can message, work, and control wheelchairs and smart homes.
- **Restored speech:** real-time voice for people who have lost the ability to speak.
- **Neurological treatment:** closed-loop stimulation for epilepsy, depression and movement disorders.
- **Longer term:** sensory restoration (vision, touch) and, more speculatively, augmentation.

Future scenarios

2027
First BCIs pursue full market approval for communication in paralysis.

2030
Reliable, home-use speech restoration at near-conversational speed for ALS.

2035
BCIs as an established medical device category with insurance coverage.

2040+
Bidirectional interfaces for sensation; vision restoration possibly maturing.

- **Optimistic:** safe, durable implants become routine treatment for paralysis and speech loss.
- **Base case:** steady clinical progress for severe disability; consumer uses remain far off.
- **Pessimistic:** durability or safety problems slow approvals for years.

Effects and risks

- **Surgery and durability:** implants must work for years; tissue reactions and thread movement degrade signals.
- **Neural data privacy:** brain data reveals intentions and health; it needs the strongest protection.
- **Hype gap:** company announcements often run ahead of peer-reviewed results.
- **Access and cost:** who pays, and who maintains devices if a company fails?

What to watch

- Peer-reviewed results from Neuralink's speech and motor trials.

- A registered Blindsight human trial.
- Full market approvals and insurance coverage decisions.
- Long-term (5+ year) implant reliability data.

My take

A BCI is a production system inside a human being. I'd judge them like critical infrastructure: uptime over years, safe failure modes, remote monitoring, secure updates and strict data protection. The breakthroughs are real; what turns them into medicine is boring, rigorous reliability engineering.

Key players & programs

Neuralink	N1 implant, R1 robot; 21+ participants by early 2026; speech and Blindsight efforts [1][2]
UC Davis / academic consortia	Speech BCI with ~97% word accuracy (NEJM 2024) [3]
Precision Neuroscience	Layer 7 thin-film interface; FDA 510(k) in April 2025 [4]
Synchron	Endovascular Stentrode, no open-brain surgery
Blackrock Neurotech / BrainGate	Utah-array research since the 2000s
Paradromics	High-bandwidth Connexus implant, early human studies

Frequently asked questions

What is a brain-machine interface?
A system that records brain activity and translates it into commands, such as moving a cursor, typing or producing synthesized speech.

How many people have a Neuralink implant?
Neuralink reported 21 participants in early 2026, with trials expanding to the UK, Canada and the UAE.

Can Neuralink restore vision?
Its Blindsight device has FDA Breakthrough Device designation, but as of late 2026 no publicly registered human trial had shown vision restoration.

Which BCI has the best published speech results?
An academic speech BCI from UC Davis decoded attempted speech with about 97% word accuracy for a man with ALS, published in NEJM in 2024.

Glossary

Utah array: A widely used grid of penetrating microelectrodes for recording neurons.

Endovascular: Placed inside a blood vessel.

510(k) clearance: A US FDA pathway clearing a device that is substantially equivalent to an existing one.

Breakthrough Device designation: An FDA program that speeds review for devices treating serious conditions.

References

1. TeslaNorth — How 21 people are using Neuralink in 2026
https://teslanorth.com/2026/01/28/from-paralysis-to-neuroscience-how-21-people-are-using-neuralink-in-2026

2. MobiHealthNews — Elon Musk provides update on Blindsight implant trials
https://www.mobihealthnews.com/news/elon-musk-provides-update-blindsight-implant-trials
3. Card et al. — An accurate and rapidly calibrating speech neuroprosthesis (NEJM, 2024)
https://www.nejm.org/doi/full/10.1056/NEJMoa2314132
4. NPR — Click, speak, move: these brain implants are poised to help people with disabilities
https://wbhm.org/npr-story/click-speak-move-these-brain-implants-are-poised-to-help-people-with-disabilities
5. Paradromics — BCI benchmarking
https://www.paradromics.com/blog/bci-benchmarking

10 · BIOTECH

CRISPR and intelligence enhancement: science, limits and ethics

A deep dive into whether gene editing could make people smarter: how CRISPR, base and prime editing work, what genetics really says about intelligence, why embryo editing for IQ is not feasible today, and the ethical lines involved.

2012	Dec 2023	Thousands	2018
CRISPR-Cas9 editing described	first CRISPR therapy approved (FDA)	of variants linked to cognitive traits	first gene-edited babies (condemned)

Key takeaways

- CRISPR is already medicine: Casgevy, approved by the FDA in December 2023, treats sickle cell disease, and personalized base editing saved an infant in 2025.
- Intelligence is highly polygenic: thousands of variants each have tiny effects, and environment matters enormously.
- Editing embryos to raise intelligence is not feasible today: too many sites, unclear causal variants, pleiotropy, and safety risks.
- Polygenic embryo screening (choosing, not editing) is already sold commercially and is drawing growing ethical criticism; modeled gains are small.
- The responsible near-term goal is treating disease, including brain disease, not designing 'smarter' children.

How gene editing works

- **CRISPR-Cas9 (2012):** a guide RNA leads the Cas9 enzyme to a matching DNA sequence, where it cuts both strands; the cell's repair process then disrupts or rewrites the gene. Jennifer Doudna and Emmanuelle Charpentier received the 2020 Nobel Prize in Chemistry for it.
- **Base editing (2016):** changes a single DNA letter without a double-strand break, reducing unintended damage.
- **Prime editing (2019):** "search and replace" editing that can make small insertions, deletions and substitutions.
- **Somatic vs. germline:** somatic edits affect only the treated person; germline edits (in embryos, eggs or sperm) are inherited by future generations.

CRISPR as medicine today

- **Casgevy:** the FDA approved the first CRISPR-based therapy for sickle cell disease in December 2023 [1].
- **Personalized editing:** in 2025, an infant with a rare metabolic disorder received a base-editing therapy designed for his specific mutation, published in NEJM.
- **Pipeline:** trials target inherited blindness, high cholesterol, rare liver diseases and more.

What genetics says about intelligence

- **Heritability:** twin studies estimate that roughly half of the variation in adult cognitive test scores is associated with genetic differences, within the populations studied, but heritability is not destiny and doesn't describe individuals.
- **Polygenicity:** large genome-wide studies have linked thousands of genetic variants to cognitive ability and educational attainment, each with a tiny effect.
- **Prediction is weak:** polygenic scores explain only a modest share of differences between people, and much less within families.
- **Pleiotropy:** the same variants affect other traits, including some psychiatric risks; selecting for one can change others [2].

Why editing for IQ isn't feasible

1. **Too many targets:** meaningful change would need edits at hundreds or thousands of sites; today's tools are tested on one or a few.
2. **Unknown causality:** most associated variants are markers near the true causal ones.
3. **Safety:** each edit risks off-target changes and mosaicism (only some cells edited), a serious risk in an embryo.
4. **Context:** effects depend on ancestry, environment and other genes; results from one population don't transfer cleanly.

A 2025 Nature analysis modeled heritable polygenic editing and found large *theoretical* effects on disease risk, while stressing the approach is speculative and raises eugenics and inequality concerns [3].

What is happening instead

- **Polygenic embryo screening:** during IVF, companies score embryos' genomes and parents choose among them. Modeling suggests the expected gain in cognitive scores is small, and critics argue it can produce eugenic outcomes with less regulatory attention than editing [2][4].
- **Germline editing:** prohibited or heavily restricted in most countries since He Jiankui's 2018 gene-edited twins, which were widely condemned and led to his imprisonment. In 2026, commentary noted falling technical barriers and new startups, with ethical scrutiny rising in response [4].

How it could change life

Responsibly, the biggest impact on human minds will come from **treating brain disease**: inherited neurodevelopmental disorders, and lowering risk for conditions such as Alzheimer's. Gene therapies that keep brains healthy for longer are far more realistic than gene edits that make healthy brains "smarter".

Future scenarios

2027–2030	2030s	2040	Long term
More somatic CRISPR, base and prime editing therapies approved; costs start falling.	Regulation of embryo screening tightens; international debate on any germline use for disease.	Possible narrowly defined germline use to prevent severe single-gene disease, if ever permitted.	Cognitive enhancement by editing remains scientifically uncertain and ethically contested.

- **Optimistic:** gene editing cures many inherited diseases affordably worldwide.

- **Base case:** steady therapeutic progress; enhancement stays off-limits.
- **Pessimistic:** unregulated "enhancement" markets emerge in permissive jurisdictions, deepening inequality.

Ethics

- **Consent:** future people can't consent to germline changes.
- **Eugenics:** history shows how "improving" populations leads to discrimination and harm.
- **Inequality:** enhancement available only to the wealthy could entrench advantage across generations.
- **Diversity:** traits linked to cognitive variation also connect to creativity and neurodiversity; "optimizing" is not neutral.

What to watch

- New somatic gene-editing approvals and prices.
- Regulation of polygenic embryo screening.
- International summits and national laws on heritable editing.
- Peer-reviewed evidence (not marketing) from embryo-screening companies.

My take

I'm excited about CRISPR as medicine and skeptical of it as an IQ upgrade. Intelligence grows from health, nutrition, education, relationships and good tools. If we want smarter societies, the highest-return investments are still schools, healthcare, and AI that helps everyone learn and think.

Key players & programs

Casgevy (Vertex / CRISPR Therapeutics)	First approved CRISPR therapy, sickle cell disease (2023) [1]
Embryo-screening companies	Polygenic screening during IVF; ethically contested [2]
Germline-editing startups (2025–26)	Under rising ethical scrutiny [4]
International summits and WHO	Governance frameworks for human genome editing

Frequently asked questions

Can CRISPR make people smarter?
Not with today's science. Intelligence involves thousands of genetic variants with tiny effects plus strong environmental influence, and editing many sites in an embryo safely is not possible.

What is CRISPR used for today?
Treating disease. The FDA approved Casgevy for sickle cell disease in December 2023, and personalized base editing treated an infant with a rare metabolic disorder in 2025.

What is polygenic embryo screening?
Scoring IVF embryos' genomes for many traits and choosing among them. It is sold commercially, but modeled gains are small and it is ethically controversial.

Is editing human embryos legal?
Heritable human genome editing is prohibited or heavily restricted in most countries, especially since the 2018 gene-edited twins case.

Glossary

Polygenic: Influenced by many genes, each with a small effect.

Germline editing: Editing embryos, eggs or sperm so changes are inherited.

Pleiotropy: When one gene or variant affects several different traits.

Mosaicism: When only some of an organism's cells carry an edit.

References

1. FDA — FDA approves first gene therapies to treat patients with sickle cell disease (Dec 2023) https://www.fda.gov/news-events/press-announcements/fda-approves-first-gene-therapies-treat-patients-sickle-cell-disease
2. Genetic Literacy Project — Should we regulate or ban embryo selection based on IQ? (2025) https://geneticliteracyproject.org/2025/04/07/viewpoint-should-we-regulate-or-ban-embryo-selection-based-on-iq/
3. Visscher et al. — Heritable polygenic editing: the next frontier in genomic medicine? (Nature, 2025) https://pmc.ncbi.nlm.nih.gov/articles/PMC11735401
4. Genetic Literacy Project — With technical barriers to human embryo editing falling, ethical concerns escalate (Nov 2026) https://geneticliteracyproject.org/2026/11/05/with-technical-barriers-to-human-embryo-editing-falling-ethical-concerns-escalate/
5. Milbank Quarterly — Much ado about CRISPR germline editing https://www.milbank.org/quarterly/opinions/much-ado-about-crispr-germline-editing/

Online: https://mohsinkhanshaik.com/research/crispr-intelligence · Facts are cited; outlooks and "My take" are the author's opinion.

11 · QUANTUM

Quantum AI: what’s real, what’s next

A deep dive into where quantum computing and AI meet: Google's verifiable 'Quantum Echoes' advantage, why quantum machine learning hasn't beaten GPUs, how AI is already improving quantum computers, and the realistic road to useful machines.

13,000×	105	99.88%	2033
Quantum Echoes speed-up claim	qubits on Google Willow	two-qubit gate fidelity	DARPA QBI utility-scale target

Key takeaways

- 'Quantum AI' means two things: quantum computers running AI, and AI improving quantum computers. The second is happening now; the first is mostly research.
- In October 2025 Google reported Quantum Echoes on its 105-qubit Willow chip: about 13,000 times faster than the best classical method, and verifiable, published in Nature.
- Broad quantum speed-ups for machine learning have not been shown; several claimed advantages were matched by better classical algorithms.
- AI is already essential to quantum hardware: DeepMind's AlphaQubit decoder (Nature, 2024) beat standard error-correction decoders on Google's data.
- Early useful applications are most likely in simulating molecules and materials, not chatbots.

Two meanings of quantum AI

- **Quantum for AI:** algorithms that use qubits to speed up learning, optimization or sampling.
- **AI for quantum:** machine learning that designs circuits, calibrates hardware, decodes errors and interprets results.

Where quantum hardware stands

- **Willow (Dec 2024):** Google showed error rates falling as error-correcting codes grew, a key threshold for scaling.
- **Quantum Echoes (Oct 2025):** an out-of-time-order-correlator algorithm on 65 of Willow's 105 qubits ran about 13,000 times faster than the best classical approach on a top supercomputer, and its expectation values can be checked on other quantum computers, making it the first verifiable quantum advantage [1][2].
- **Fidelities:** 99.97% single-qubit gates, 99.88% entangling gates and 99.5% readout [1].
- **Chemistry link:** a companion experiment matched NMR measurements on real molecules [2][4].
- **Independent testing:** DARPA's QBI advanced four companies to Stage C in October 2026 to judge whether utility-scale machines are possible by 2033 [3].

Quantum machine learning: the honest picture

- **Data loading:** getting large classical datasets into quantum states can cancel out speed-ups.

- **Dequantization:** since 2018, several proposed exponential speed-ups (for example in recommendation systems) were matched by new classical algorithms.
- **Barren plateaus:** many parameterized quantum circuits become impossible to train as they grow because gradients vanish.
- **Where hope remains:** problems that are naturally quantum (molecules, materials, quantum sensor data) and certain sampling tasks.

AI for quantum: already working

- **Error decoding:** DeepMind and Google's AlphaQubit (Nature, 2024) used a neural network to identify errors more accurately than standard decoders on surface codes; follow-up work tackles its speed for real-time use [5].
- **Calibration and control:** machine learning tunes thousands of hardware parameters faster than humans.
- **Algorithm discovery:** AI systems search for better circuits and compilation strategies.

How it could change life

- **Medicine and materials:** accurate simulation of molecules for drugs, catalysts, batteries and fertilizers.
- **Science:** new tools for physics and chemistry that classical computers can't simulate.
- **Security:** the same progress makes post-quantum migration urgent.
- **Not soon:** quantum chatbots or quantum-trained large language models.

Future scenarios

2027–2028	2029–2030	2033	2035+
More verifiable advantages in physics and chemistry; AI decoders standard.	First machines with roughly a hundred or more logical qubits on vendor roadmaps.	QBI's judgment on utility-scale feasibility.	Routine quantum simulation in R&D, if error correction scales.

- **Optimistic:** error-corrected machines deliver chemistry breakthroughs by the early 2030s.
- **Base case:** useful simulation for specific problems arrives in the 2030s; QML stays niche.
- **Pessimistic:** scaling stalls on error rates, wiring and cost.

Effects and risks

- **Hype:** "quantum AI" marketing often outruns evidence; demand verifiable benchmarks.
- **Security:** shorter timelines to breaking RSA and ECC.
- **Concentration:** few organizations can build these machines.

What to watch

- Logical-qubit counts with error rates below physical ones.
- New verifiable-advantage results tied to useful problems.
- QBI Stage C evaluations.
- Any QML result that beats strong classical baselines.

My take

I'm optimistic about quantum for chemistry and skeptical of near-term quantum machine learning. The most useful thing engineers can do today is make systems crypto-agile and keep an eye on verifiable, independently benchmarked results.

Key players & programs

Google Quantum AI	Willow (2024), Quantum Echoes verifiable advantage (2025) [1][2]
Google DeepMind	AlphaQubit neural error decoder (Nature, 2024) [5]
DARPA QBI	Independent benchmarking; Stage C: Atom Computing, Diraq, IBM, IonQ [3]

Frequently asked questions

What is quantum AI?
Either using quantum computers to run or speed up AI, or using AI to build and improve quantum computers. The second is already common; the first remains mostly research.

What is Google's Quantum Echoes?
An algorithm run on Google's 105-qubit Willow chip in 2025 that was about 13,000 times faster than the best classical method and verifiable on other quantum computers, published in Nature.

Will quantum computers make ChatGPT-style AI faster?
Not in the foreseeable future. Broad quantum speed-ups for machine learning haven't been demonstrated, and loading large data into qubits is a bottleneck.

How does AI help quantum computers?
AI decodes errors (for example DeepMind's AlphaQubit), calibrates hardware and helps design better circuits.

Glossary

Qubit: The basic unit of quantum information; it can be in a superposition of 0 and 1.

Logical qubit: A qubit encoded across many physical qubits using error correction.

Dequantization: Finding a classical algorithm that matches a claimed quantum speed-up.

Verifiable advantage: A quantum computation that beats classical methods and whose result can be independently checked.

References

1. Google — Our quantum hardware: the engine for verifiable quantum advantage
   https://blog.google/technology/research/quantum-hardware-verifiable-advantage/
2. Google Research — A verifiable quantum advantage
   https://research.google/blog/a-verifiable-quantum-advantage/
3. DARPA — Quantum Benchmarking Initiative
   https://www.darpa.mil/research/programs/quantum-benchmarking-initiative
4. Science News — Quantum 'echoes' reveal the potential of Google's quantum computer
   https://www.sciencenews.org/article/quantum-echoes-google-computer
5. arXiv — Learning to decode in parallel: neural network for real-time QEC (2026, on AlphaQubit)
   https://arxiv.org/pdf/2601.09921

12 · BRAIN & NEUROTECH

Princeton's 3D maps of the mind: FlyWire and MICrONS

A deep dive into connectomics: how Princeton-led teams mapped every neuron in a fruit fly brain and a cubic millimeter of mouse cortex in 3D, how AI and gamers made it possible, and what wiring diagrams mean for medicine and AI.

139,255	54.5 million	524 million	1.6 PB
neurons in the fly connectome	fly synapses	synapses in 1 mm^3 of mouse cortex	MICrONS data release

Key takeaways

- A connectome is a complete 3D wiring diagram of a nervous system: every neuron and every connection.
- FlyWire (Nature, October 2024), led by Princeton, mapped all 139,255 neurons and about 54.5 million synapses of an adult fruit fly brain.
- MICrONS (April 2025), co-led by Princeton, Baylor and the Allen Institute, mapped one cubic millimeter of mouse visual cortex: over 524 million synapses and 200,000+ cells, plus recordings of the same neurons' activity.
- AI did most of the tracing; humans, including volunteer gamers, proofread, a model of human–AI collaboration at scale.
- Connectomes are reference maps, like the genome, for understanding brain function, disease and more efficient AI.

A short history

- **1986:** the 302-neuron connectome of the worm *C. elegans* was published after years of manual tracing.
- **2012:** Sebastian Seung's book *Connectome* argued "you are your connectome", and his lab launched Eyewire, a game where volunteers trace neurons.
- **2020:** the "hemibrain" mapped about 25,000 neurons in part of a fly brain [3].
- **2024:** FlyWire completed the whole adult fly brain [3][6].
- **2025:** MICrONS released the largest mammalian wiring diagram with matching activity data [1][4].

How a connectome is made

1. **Prepare and slice:** a small piece of brain is stained and cut into thousands of ultra-thin sections, or milled layer by layer.
2. **Image:** electron microscopes capture each slice at nanometer resolution, producing petabytes of images.
3. **Align:** software stitches slices into a 3D volume.
4. **Segment:** convolutional neural networks trace every neuron and detect synapses.
5. **Proofread:** scientists and volunteers fix merge and split errors.

6. **Analyze:** the graph of neurons and synapses is studied like a network.

FlyWire: the whole fly brain

- Published in Nature on 2 October 2024 as a package of papers led by Princeton with partners including Cambridge [2][3].
- Covers 139,255 neurons and about 54.5 million chemical synapses in an adult female fruit fly [3][6].
- Proofreading took an estimated 33 person-years from the FlyWire Consortium; Seung said that without AI the project would have needed almost 50,000 person-years [2][3].
- Data is openly explorable online through the FlyWire Codex [6].

MICrONS: a cubic millimeter of mouse cortex

- One cubic millimeter of mouse visual cortex reconstructed, with more than 524 million synapses, over 200,000 cells and about 4 km of axons [4][5].
- Before slicing, the team recorded how neurons responded to what the mouse saw, so structure and function are linked [5].
- A 1.6-petabyte release published in Nature and Nature Methods on 9 April 2025, co-led by Princeton, Baylor College of Medicine and the Allen Institute [1][4][5].

How it could change life

- **Medicine:** comparing healthy and disease-model circuits can reveal what goes wrong in epilepsy, autism or neurodegeneration.
- **Neuroscience:** testing theories of how circuits compute against real wiring.
- **AI:** real connectivity patterns inspire more efficient neural network architectures.
- **Neurotech:** ground truth for brain-computer interfaces and organoid computing.

Future scenarios

2027–2028	2030	2035	2040+
Larger mouse-brain regions; disease-model comparisons.	Whole-mouse-brain connectome efforts become feasible as AI tracing improves.	Detailed maps of human brain regions linked to function and disease.	Connectome-constrained brain simulations used in research and drug discovery.

- **Optimistic:** a whole-mouse-brain connectome by the early 2030s transforms neuroscience.
- **Base case:** steadily larger regional maps with linked activity data.
- **Pessimistic:** exabyte-scale storage and proofreading costs slow progress.

Effects and risks

- **Data scale:** a whole mouse brain would be exabytes; storage, compute and analysis become core challenges.
- **Interpretation:** a wiring diagram isn't an explanation; activity, chemistry and models are also needed.
- **Ethics:** human connectomics raises questions about privacy, identity and consent.

What to watch

- Whole-mouse-brain connectome milestones.
- Automated proofreading that reduces human effort.
- Studies linking connectome differences to disease.

Connection to my work

Connectomics is one of science's great data-infrastructure feats: petabyte pipelines, GPU-heavy segmentation, global collaboration tools and open data portals. Reliability, cost and observability at that scale are exactly the problems I work on, and they will define how fast neuroscience can go next.

My take

FlyWire and MICrONS show what happens when AI pipelines and human judgment work together at scale. They're also a reminder that the hardest part of big science is often the infrastructure underneath it.

Key players & programs

Princeton (Seung lab)	Co-led FlyWire and MICrONS; Eyewire [1][2]
Allen Institute and Baylor College of Medicine	MICrONS co-leads [4]
University of Cambridge and FlyWire Consortium	FlyWire partners and proofreading community [3]
Zetta AI	Princeton spin-off doing connectome alignment and proofreading

Frequently asked questions

What is a connectome?
A complete wiring diagram of a nervous system, showing every neuron and the connections between them, usually reconstructed in 3D from electron-microscope images.

What did Princeton's FlyWire map?
Every one of the 139,255 neurons and about 54.5 million synapses in an adult fruit fly brain, published in Nature in October 2024.

What is MICrONS?
A 2025 project co-led by Princeton, Baylor and the Allen Institute that mapped one cubic millimeter of mouse visual cortex, over 524 million synapses and 200,000+ cells, together with recordings of the neurons' activity.

How did AI help map the brain?
Neural networks traced neurons and detected synapses automatically; humans proofread. Without AI, FlyWire would have taken almost 50,000 person-years.

Glossary

Synapse: A connection where one neuron passes a signal to another.

Segmentation: Automatically outlining each neuron in microscope images.

Proofreading: Human correction of errors in automated reconstructions.

Electron microscopy: Imaging with electrons to see structures at nanometer scale.

References

1. Princeton Engineering — Scientists map vision centers of the mammalian brain (Apr 2025)
   https://engineering.princeton.edu/news/2025/04/09/scientists-map-vision-centers-mammalian-brain
2. AAU — Mapping the entire fly brain, a step toward understanding
   https://www.aau.edu/newsroom/mapping-entire-fly-brain-step-toward-understanding
3. University of Cambridge — First map of every neuron in an adult fly brain complete
   https://www.cam.ac.uk/research/news/first-map-of-every-neuron-in-an-adult-fly-brain-complete
4. Allen Institute — Scientists complete largest wiring diagram and functional map of the brain to date
   https://alleninstitute.org/news/scientists-complete-largest-wiring-diagram-and-functional-map-of-the-brain-to-date
5. NIH — Scientists map unprecedented detail of connections and visual perception in the mouse brain
   https://www.nih.gov/news-events/news-releases/scientists-map-unprecedented-detail-connections-visual-perception-mouse-brain
6. Dorkenwald et al. — Neuronal wiring diagram of an adult brain (Nature, 2024)
   https://pubmed.ncbi.nlm.nih.gov/39358518/