BRIEFING 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.

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

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.