


Photonic Computing Takes a Step Toward Fruition

Two newly developed computer chips, powered in part by light, have tackled complex computing tasks once considered out of reach for photonic systems.

By **Ananya Palivela**

Photonic computing uses light instead of electrons to store information and to perform computations, offering the promise of faster analog computer processing with lower energy overhead. Now two companies have demonstrated photonic processors that take a step toward this reality. One device, built by researchers at Lightintelligence, a photonics computing company in Singapore, solves hard optimization problems faster than traditional processors [1]. The other, developed by Lightmatter, a photonics computing company in California, runs modern artificial intelligence (AI) models with high accuracy and high energy efficiency [2].

Light signals don't interfere with each other like electrical signals do, so photonics-based systems can process many computations simultaneously. However, analog photonic computers are noisy, and researchers have yet to integrate them into standard computing systems despite a decade of trying.



A quadcore photonic processor developed by Lightmatter.
Credit: Lightmatter

The Lightintelligence team, led by Bo Peng, built a chip called Photonic Arithmetic Computing Engine (PACE), combining photonic components with electronic circuits. The chip was designed to solve optimization problems by mimicking a physical model known as the Ising spin model. In the Ising model, variables (or “spins”) are arranged on a graph and interact with one another in a way that favors low-energy configurations of the spins.

Many real-world problems, such as image matching and network partitioning, can be framed as finding the lowest-energy state of such a system. The PACE chip performs a matrix-vector multiplication in a feedback loop to find these states. It multiplies a given variable configuration by a matrix representing the problem (encoded in light), electronically processes the result, and feeds the answer back to the photonic processor to update the variables. The system converges toward a solution by repeating this loop.

Using this method, the Lightintelligence team reconstructed a noisy image of a cat and solved optimization problems 500 times faster than a modern graphics processing units (GPUs). The demonstration shows that you can reach impressively high computing speeds with the parallelism of optics, says Dirk Englund, a computer scientist at MIT who was not involved in either study. “The numbers are quite compelling, and this could be just the beginning.”

While the Lightintelligence chip is designed for solving optimization problems, Lightmatter’s system tackles a broader range of tasks. The company’s photonic system can run standard AI models such as ResNet and BERT using photonic

tensor cores—specialized computing units that use light instead of electricity to perform matrix–vector multiplications. These cores are then controlled by electronic interface chips. “We show that you can build a computer that’s not based on transistors and run state-of-the-art workloads,” says Nicholas Harris, cofounder and CEO of Lightmatter.

One of the most serious bottlenecks in current AI systems is not the speed of the computation but the delay and energy cost of moving data between processors. “If you had a processor that used zero energy and was infinitely fast, you wouldn’t even get a 2 times speedup because the machine is mostly waiting for data,” Harris says. So, they designed a digital interface chip to go along with the photonic tensor cores, enabling fast and efficient communication between components.

Small errors in light intensity or signal drift can quickly throw off complex calculations, especially in AI tasks that depend on numerical accuracy. Lightmatter addressed this potential problem by using active calibration, where the chip adjusts its power if the light intensity changes. The team also used a specialized number format called adaptive block floating point. AI models rely on precise numbers to make accurate predictions. In traditional formats, each number can have its own exponent, allowing for high precision. But in photonic hardware, assigning and managing individual exponents is costly and difficult. Lightmatter’s format groups numbers to share a common exponent, simplifying the hardware and helping reduce noise. “This adaptive block floating point method is a clever mathematical solution to a challenge that’s held the field back,” Englund says.

The Lightmatter system can perform image classification,

natural language processing, and reinforcement learning tasks. “We have the photonic computer playing Pac-Man!” Harris says. The models did not need to be retrained or adapted to the hardware. “If you have a desktop PC at home, you can plug this processor into your computer.”

The chip achieved a 7- to 10-bit precision, is at least 10 times more energy-efficient than traditional GPUs, and delivers comparable speed on standard AI tasks. “It was designed to run at 400 trillion operations per second,” Harris says. “But we achieved around 80 trillion. Getting to 400 [trillion] is something I think is well within reach.”

Photonic computing continues to face challenges like analog noise, integration difficulties, and the high energy cost of moving data between chips. Englund says that these two results show that photonic computing is finally stepping out of the lab and into real-world relevance, and photonics can tackle vastly different problems with architectures tailored to each. These two new demonstrations show that photonic processors can now scale up, operate reliably, and perform real-world tasks. “It’s an incredible leap ahead,” Englund says. “I was looking at a paper we published eight years ago, and it’s staggering how much progress there has been.”

Ananya Palivela is a freelance science writer based in Dallas, Texas.

REFERENCES

1. S. Hua *et al.*, “An integrated large-scale photonic accelerator with ultralow latency,” *Nature* **640**, 361 (2025).
2. S. R. Ahmed *et al.*, “Universal photonic artificial intelligence acceleration,” *Nature* **640**, 368 (2025).