Vision and reality: artificial intelligence still has a long way to go before it truly understands complex relationships in physics. But today's algorithms are already capable of inspiring researchers and suggesting surprising designs for experiments.

IMAGE: ADOBESTOCK

ARTIFICIAL INSPIRATION

TEXT: ROLAND WENGENMAYR

Whether making a medical diagnosis, searching for materials for the energy revolution, or predicting protein structures, artificial intelligence algorithms are an effective tool in many scientific fields today. But are they useful in physics, where the goal is to understand the fundamental processes of nature? Mario Krenn and Florian Marquardt are already assisted by AI at the Max Planck Institute for the Science of Light, where they are getting a feel for what the algorithms can and cannot (yet) do.

Artificial intelligence, or AI for short, is booming. Many people use ChatGPT, for instance. The most famous example of an AI based on a large language model, ChatGPT helps people conduct research or write texts. AI can also be used to generate images or videos from text prompts and has long served as a tool in the art world. But what is its role in the sciences?

AI is already well established in the life sciences and chemistry. The AlphaFold program from DeepMind gained prominence in biology due to its ability to calculate protein structures. You may recall that DeepMind caused a furor with its program AlphaGo in 2016, when it beat the world's strongest Go player, Lee Sedol. It was a sensation, because there are so many possible next moves in Go that no computer can calculate them all. For that reason, AlphaGo had to learn by training like a human, developing a feeling for patterns in the combinations of stones played on the board – and hence a kind of understanding of what constitutes a smart move. In the process, it definitely benefited from brute-force calculating power. It was able to train against itself, so to speak, by playing millions of games, whereas professional Go players only ever play a few thousand.

Despite the widespread application of AI since then, most programs function as a black box, meaning that a user gets a useful result without understanding how it was arrived at. That's often sufficient – for example, when the user is searching for a protein with a specific function. In this case, what matters is to understand why the structure discovered by the AI does what it should. But in physics, the most fundamental of all sciences, a black box undermines the goal of understanding a physical system. Researchers increasingly turn to AI in physics as well, but mostly for applications where the use of a black box doesn't inhibit understanding. However, experts are debating whether AI could become efficient enough to un-

⟶

derstand complex physical systems even better than humans. If so, and if they were able to explain things to their human colleagues, would they gain an equal footing as artificial physicists? And could this inspire new ideas in physics?

Eureka moment in machine learning

Mario Krenn and other physicists consider this type of AI an artificial muse. They wrote an article about it in the December 2022 issue of the trade journal Nature Reviews Physics. We are sitting in the cafeteria of the Max Planck Institute for the Science of Light in Erlangen. With us is the Director of the Theory Department, Florian Marquardt. Like Krenn, he has employed several methods for machine learning in the past few years and continues to refine them with his team. Krenn earned his Ph.D. in experimental quantum optics in Vienna under Anton Zeilinger, winner of the Nobel Prize for Physics in 2022. He changed course radically following a eureka moment involving machine learning. “I never set foot in a lab again,” he says with a laugh. He is considered a pioneer in the use of AI in physics. Today he heads a research group at the Institute called the Artificial Scientist Lab; the name alone conveys a vision of an artificial physicist.

We’ll say more about what motivated Krenn to change course in 2014, but first, it is necessary to explain the conditions under which an AI could be considered equal to a physicist. “To begin with, we have to understand how human researchers work,” Krenn

AI can contribute to a scientific understanding of physics at several difficulty levels: as a kind of virtual microscope, as a source of inspiration, and as a participant with an understanding of its own.

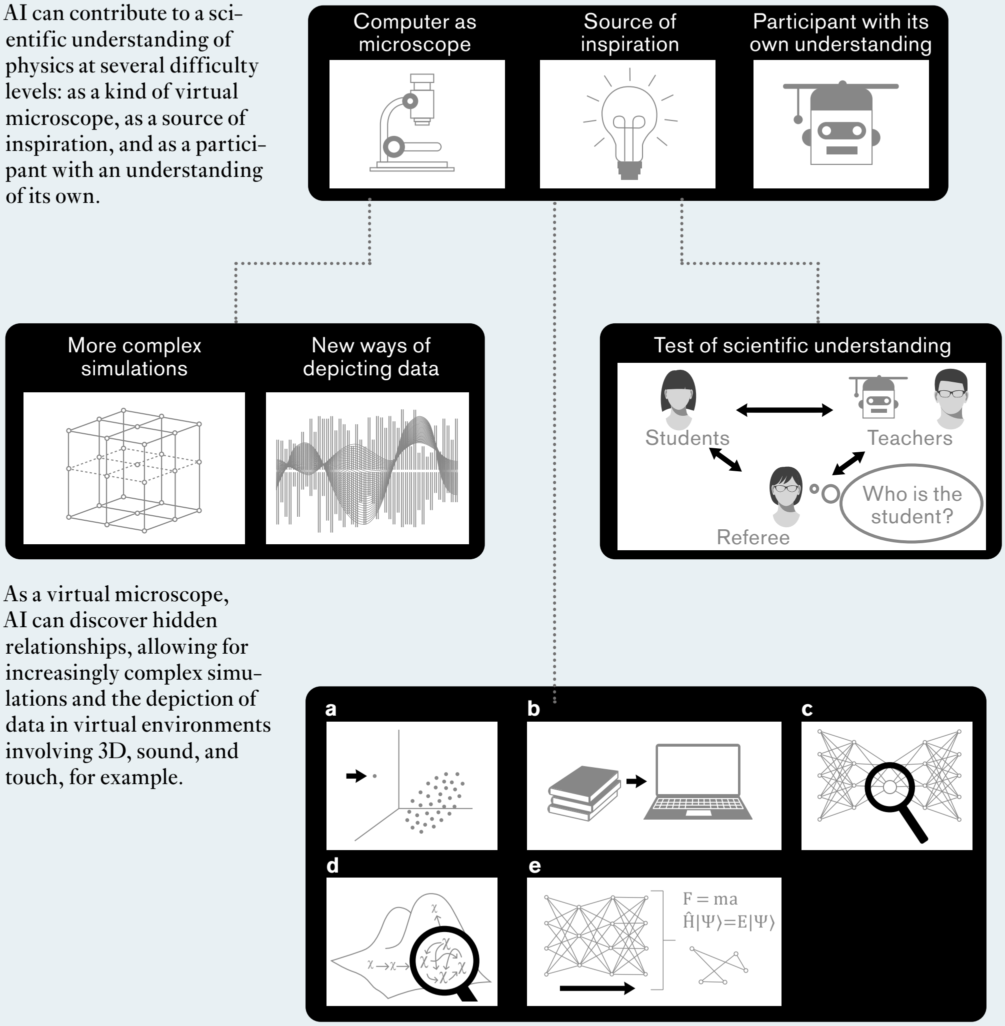

As a virtual microscope, AI can discover hidden relationships, allowing for increasingly complex simulations and the depiction of data in virtual environments involving 3D, sound, and touch, for example.

At the third and still unattained level, AI could achieve a physical understanding on its own. This could be proven with a test, in which an AI or a human explains a scientific theory to students. If a referee is no longer able to distinguish between the level of knowledge possessed by the teacher and the student, the theory has been understood and communicated by the teacher – whether AI or human.

It can draw inspiration from surprises in the datasets (a) or the scientific literature (b). Moreover, it can discover new and unexpected concepts by investigating scientific models (c), exploring a data space with programmed curiosity or creativity (d), or providing solutions in the form of, for example, mathematical formulas that researchers can interpret.

GRAPHIC: GCO BASED ON KRENN, M., POLLICE, R., GUO, S.Y. ET AL. ON SCIENTIFIC UNDERSTANDING WITH ARTIFICIAL INTELLIGENCE. NAT REV PHYS 4, 761–769 (2022).

emphasizes, “why they are creative, how they are creative, and what makes them curious.” The question is therefore what motivates people to do their research. “Once we understand that, we have a better chance of doing truly autonomous, automated science,” he says. Marquardt agrees, adding: “At the same time, we’re learning something about what we humans are doing in science. It’s far from clear that all of our priorities in research are really so objective!”

Consequently, we can establish as a starting point that an artificial physicist must be able to motivate itself to the same degree as a human. It sounds banal, but an example illustrates how challenging this vision is. “Take, for instance, a key question in solid-state physics,” Marquardt suggests. “How can I manufacture a room-temperature superconductor?” Disregarding the fact that even a question this specific and unanswered is still too open and general for today’s AI, an artificial physicist would have to arrive at this question by itself and rank it as important. The AI would therefore have to recognize on its own that lossless conduction of electrical power at room temperature is a desirable research objective. “But why do you want to transport electricity losslessly?” Krenn asks, inquiring into the next meta-level of knowledge. The AI would have to ask and answer this question itself without external prompting. In sum, it would have to know that electrical energy is crucially important to our society. But this goes into social issues that are far outside the purview of physics.

This example illustrates just how challenging the human traits of creativity and curiosity really are. AI is still a long way away. It might be closer to achieving a kind of understanding of physical theories, however. One question that arises here is what exactly it means to understand a physical relationship. The conversation with Krenn and Marquardt shed light on multiple aspects that play a key role in answering this question. To answer it, physicists need an intuitive, pictorial, or model-like representation, even if it’s only an abstract mathematical description. In the case of AlphaGo, AI has already shown itself capable of this type of intuition – in the special case of the situation on a game board. However, to understand also means to be able to transfer insights and solutions from one area to another. “When an AI becomes familiar with a concept in one context, possibly by learning about it in the scientific literature like us, then maybe it can recognize that the concept can also be applied in another context,” says Marquardt. Ultimately, an AI would have to be capable of explaining a relationship to humans as well, possibly with the help of a language model. Krenn and Marquardt believe an AI can manage it. But first, AI has a lot to learn.

In many specialized tasks, AI already leaves human physicists in the dust. And that’s exactly what Krenn and Marquardt are counting on in their research. One method they rely on is artificial neural networks. The networks simulate interconnected nerve cells, which learn by strengthening or weakening certain neural connections through training. “Artificial neural networks are only one method, however. The spectrum of AI is much broader,” Marquardt points out. “But all AI methods have in common the fact that they help deal with complexity.” This includes being able to discover hidden patterns and solve mathematical optimization problems, says Marquardt. Through training with millions of images and other methods, AI learns to identify objects such as “cars” or “eagles” in a highly diverse range of perspectives and situations.

Solutions for quantum error correction

It is precisely this ability to recognize patterns that Marquardt leverages in his work. Several years ago, one of his teams trained an AI to find solutions for quantum error correction. The next generation of quantum computers will rely on this kind of correction, because their highly sensitive quantum bits are necessarily subject to disturbances from the environment. One of the peculiarities of the quantum world is that during a quantum calculation it isn’t possible to take measurements to determine whether the qubits still contain the correct values. For that reason, quantum error correction has to cleverly avoid direct measurement. It’s almost as if you were playing Go against an opponent whose white stones you couldn’t see, and the only way to determine where they are is by placing your own stones carefully. As with other problems, quantum error correction therefore involves recognizing patterns. Furthermore, the Erlangen researchers’ AI-based program has detected new sequences of quantum operations for certain correction algorithms.

With the help of AI, Florian Marquardt’s group has also discovered other fault-tolerant programming methods for quantum computers, as well as designs for photonic integrated circuits – the optical counterparts of electronic integrated circuits. In addition, his group has developed approaches for neuromorphic computer architectures. Because of how it works, AI currently requires a lot of energy.

⟶

SUMMARY

Artificial intelligence is exceptionally good at detecting patterns in large quantities of data and reducing the complexity of relationships. This enables it to design experiments in the fields of, for example, quantum physics or gravitational wave physics.

In the future, AI might also be capable of understanding physical relationships when given a model of them, translating concepts from one field to another, and explaining a relationship to humans.

To be on a par with human scientists, AI would have to be capable of independently deriving questions based on the needs of society. That goal is still a long way off.

Neuromorphic chips would be much more sustainable. They were inspired by our brains, which only use enough electricity to power a 20-Watt bulb.

Krenn's enlightening experience in 2014 was likewise guided by AI. At the time, his team led by Anton Zeilinger wanted to generate an especially complex form of entanglement between light quanta, photons. Entanglement is a central tool of quantum information technology. Roughly speaking, the quantum states of individual quantum objects, such as photons, are overlapped so that together they form a large quantum system. It's a bit like a rowing crew whose members are so highly synchronized that they row like a single super-athlete.

AI designs a quantum experiment

It was unclear which experimental setup would generate the special entanglement between photons most effectively. To determine this, Krenn developed a program called Melvin that simulated all the necessary optical components, including lasers, lenses, mirrors, and detectors. As such, it was able to test millions of combinations in a brief time until it found experiments that generated the desired entanglement. Because Melvin learned which combinations were useful, the program was able to create within hours something that four physicists – three experimenters and one theoretician – had labored in vain for three months to create; it provided a working design for the experiment.

Creators of an artificial colleague: the team led by Mario Krenn (center) develops artificial intelligence that puts forward surprising solutions for experiments and advances our understanding of physics.

PHOTO: STEFAN SPANGENBERG

After this "aha" moment, Krenn dedicated himself fully to developing AI that generates suggestions for physical experiments. There was one important insight that aided him in this task: "We noticed by accident that these quantum optics experiments can be abstracted to a large extent." In fact, they can be represented as a network of mathematical graphs made of lines, known as edges and nodes. Two nodes, for example, represent two photons, while a line between them represents their entanglement. "In this abstract space, it is much easier to search for, say, new quantum experiments," Krenn explains enthusiastically. Above all, it allows a user to find the optimum solution with a minimum of nodes and edges, which can then be implemented in reality in an especially economical design with the fewest possible components. However, Krenn doesn't use artificial neural networks in his AI programs. Such networks would have to be trained with existing experimental designs, which would hardly lead to fundamentally new ideas. "We employ what are known as exploration algorithms,"

GRAPHIC: GCO BASED ON MARIO KRENN/MPI FOR THE SCIENCE OF LIGHT

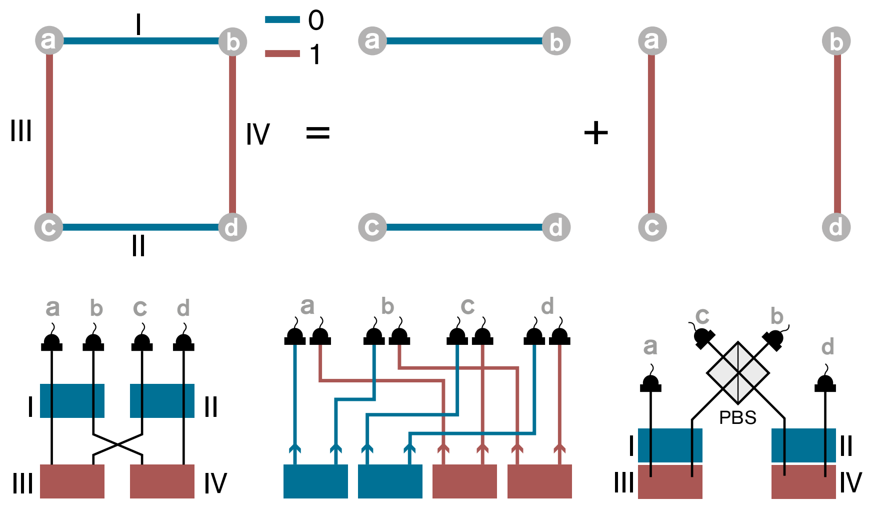

Radical simplification: a quantum experiment (below) can be depicted as a network of graphs (above). The experiment should entangle four photons, a through d (nodes in the network above). Colored lines I through IV depict pairings for the entanglement. The top left square represents the entanglement of all four photons. It results from the combination of the two adjacent graphs and can be realized in the three experiments depicted below them. The blue and red boxes represent light sources that generate individual photon pairs, while the black, hat-shaped icons symbolize detectors for incoming photons. PBS stands for an optical component that can split rays according to certain rules.

explains Krenn, "which explore the massive, abstract space of combinations for new solutions in a highly efficient way."

In the meantime, Krenn has made great strides with his research. In a work currently pre-published on the Arxiv server, an international team he was part of shows, for example, that AI can be used to develop new designs for gravitational wave detectors. Astonishingly, these designs were superior to the next generation of experiments planned for Ligo, a gravitational wave detector in the United States. Ligo became famous because it managed to discover gravitational waves whose existence had been postulated by Einstein 100 years earlier. The work was awarded the Nobel Prize in Physics in 2017. Today gravitational waves are an important new tool used by astrophysicists, for example, to detect black holes. A team led by Rana X. Adhikari is currently working on the design for the next generation, the Ligo Voyager. When the team found out Krenn had used AI to develop new quantum optics experiments, they asked him whether he would be willing to use his methods to search for new designs for the gravitational wave detectors. Collaboration resulted. As for the AI's detector designs, it will first be necessary to demonstrate in practice that unanticipated effects do not prevent them from realizing their theoretical advantages. When an experiment costs billions of dollars, however, people tend to be cautious with radical innovations.

A Nobel Prize for artificial intelligence?

With such good examples of artificial creativity, the question arises: are these harbingers of an artificial physicist? "We are now at the level where we can generate ideas," Krenn says with optimism. "On certain subjects, our AI systems can find totally new solutions that are already far more creative than human ideas in terms of novelty and utility!"

Marquardt is equally optimistic with regard to AI's applications, but far more reserved when it comes to the grand vision as a whole. For example, the question remains: when will an AI be in a position to formulate a real physical theory? Such a theory would have to be capable of being represented elegantly in clear mathematical formulas. It would have to link to existing physics and enable predictions for physical systems. It's a tall order, but Krenn is confident that within the next few years an AI will be able to provide the key idea in a discovery that wins a Nobel Prize. The Nobel Prize Committee might soon be confronted with the question of whether an AI or its creators can receive the highest prize in science.

←

GLOSSARY

ARTIFICIAL NEURAL NETWORKS
use standard computers to simulate nerve cells, which are interconnected and capable of learning by strengthening or weakening certain neural connections through training. This is only one AI method.

NEUROMORPHIC COMPUTERS
imitate the physical functioning of the human brain, which is far more energy efficient than the transistors in today's processors. Artificial neural networks can be simulated on neuromorphic computers using less energy, but are not equivalent to them.