



COMPUTING

Why AI could eat quantum computing's lunch

Rapid advances in applying artificial intelligence to simulations in physics and chemistry have some people questioning whether we will even need quantum computers at all.

By Edd Gent

November 7, 2024



STEPHANIE ARNETT/MIT TECHNOLOGY REVIEW | GETTY

Tech companies have been funneling billions of dollars into quantum computers for years. The hope is that they'll be a game changer for

fields as diverse as finance, drug discovery, and logistics.

Those expectations have been especially high in physics and chemistry, where the weird effects of quantum mechanics come into play. In theory, this is where quantum computers could have a huge advantage over conventional machines.

But while the field struggles with the realities of tricky quantum hardware, another challenger is making headway in some of these most promising use cases. AI is now being applied to fundamental physics, chemistry, and materials science in a way that suggests quantum computing's purported home turf might not be so safe after all.

The scale and complexity of quantum systems that can be simulated using AI is advancing rapidly, says Giuseppe Carleo, a professor of computational physics at the Swiss Federal Institute of Technology (EPFL). Last month, he coauthored a paper [published in *Science*](#) showing that neural-network-based approaches are rapidly becoming the leading technique for modeling materials with strong quantum properties. Meta also [recently unveiled](#) an AI model trained on a massive new data set of materials that has jumped to the top of a leaderboard for machine-learning approaches to material discovery.

 Introducing the Food issue, where we explore the intersection of tech and food. [Subscribe](#) to unlock access to every story.

Given the pace of recent advances, a growing number of researchers are now asking whether AI could solve a substantial chunk of the most interesting problems in chemistry and materials science before large-scale quantum computers become a reality.

“The existence of these new contenders in machine learning is a serious hit to the potential applications of quantum computers,” says Carleo “In my opinion, these companies will find out sooner or later that their investments are not justified.”

Exponential problems

The promise of quantum computers lies in their potential to carry out certain calculations much faster than conventional computers. Realizing this promise will require much larger quantum processors than we have today. The biggest devices have just crossed the thousand-qubit mark, but achieving an undeniable advantage over classical computers will likely require tens of thousands, if not millions. Once that hardware is available, though, a handful of quantum algorithms, like the encryption-cracking Shor's algorithm, have the potential to solve problems exponentially faster than classical algorithms can.

POPULAR

Meet the radio-obsessed civilian shapin Ukraine's drone defense

Charlie Metcalfe

Why OpenAI's new model is such a big c

James O'Donnell

Related Story

But for many quantum algorithms with more obvious commercial applications, like searching databases, solving optimization problems, or powering AI, the speed advantage is more modest. And last year, a



Quantum computing is taking on its biggest challenge: noise

For a while researchers thought they'd have to make do with noisy, error-prone systems, at least in the near term. That's starting to change.

paper coauthored by Microsoft's head of quantum computing, Matthias Troyer, showed that these theoretical advantages disappear if you account for the fact that quantum hardware operates orders of magnitude slower than modern computer chips. The difficulty of getting large amounts of classical data in and out of a quantum computer is also a major barrier.

So Troyer and his colleagues concluded that quantum computers should instead focus on problems in chemistry and materials science that require simulation of systems where quantum effects dominate. A computer that operates along the same quantum principles as these systems should, in theory, have a natural advantage here. In fact, this has

been a driving idea behind quantum computing ever since the renowned physicist Richard Feynman first proposed the idea.

The rules of quantum mechanics govern many things with huge practical and commercial value, like proteins, drugs, and materials. Their properties are determined by the interactions of their constituent particles, in particular their electrons—and simulating these interactions in a computer should make it possible to predict what kinds of characteristics a molecule will exhibit. This could prove invaluable for discovering things like new medicines or more efficient battery chemistries, for example.

But the intuition-defying rules of quantum mechanics—in particular, the phenomenon of entanglement, which allows the quantum states of distant particles to become intrinsically linked—can make these interactions incredibly complex. Precisely tracking them requires complicated math that gets exponentially tougher the more particles are involved. That can make simulating large quantum systems intractable on classical machines.

This is where quantum computers could shine. Because they also operate on quantum principles, they are able to represent quantum states much more efficiently than is possible on classical machines. They could also take advantage of quantum effects to speed up their calculations.

But not all quantum systems are the same. Their complexity is determined by the extent to which their particles interact, or correlate, with each other. In systems where these interactions are strong, tracking all these relationships can quickly explode the number of calculations required to model the system. But in most that are of practical interest to chemists and materials scientists, correlation is weak, says Carleo. That means their particles don't affect each other's behavior significantly, which makes the systems far simpler to model.

The upshot, says Carleo, is that quantum computers are unlikely to provide any advantage for most problems in chemistry and materials

science. Classical tools that can accurately model weakly correlated systems already exist, the most prominent being density functional theory (DFT). The insight behind DFT is that all you need to understand a system's key properties is its electron density, a measure of how its electrons are distributed in space. This makes for much simpler computation but can still provide accurate results for weakly correlated systems.



NEW: The Food [×] issue

Join us as we take a closer look at the future of food and how we can use technology to feed more people.

**SUBSCRIBE
TODAY**

Simulating large systems using these approaches requires considerable computing power. But in recent years there's been an explosion of research using DFT to generate data on chemicals, biomolecules, and materials—data that can be used to train neural networks. These AI models learn patterns in the data that allow them to predict what properties a particular chemical structure is likely to have, but they are orders of magnitude cheaper to run than conventional DFT calculations.

This has dramatically expanded the size of systems that can be modeled—to as many as 100,000 atoms at a time—and how long simulations can run, says Alexandre Tkatchenko, a physics professor at the University of Luxembourg. “It’s wonderful. You can really do most of chemistry,” he says.

Olexandr Isayev, a chemistry professor at Carnegie Mellon University, says these techniques are already being widely applied by companies in chemistry and life sciences. And for researchers, previously out of reach problems such as optimizing chemical reactions, developing new battery materials, and understanding protein binding are finally becoming tractable.

As with most AI applications, the biggest bottleneck is data, says Isayev. Meta's recently [released materials data set](#) was made up of DFT calculations on 118 million molecules. A model trained on this data achieved state-of-the-art performance, but creating the training material took vast computing resources, well beyond what's accessible to most research teams. That means fulfilling the full promise of this approach will require massive investment.

Modeling a weakly correlated system using DFT is not an exponentially scaling problem, though. This suggests that with more data and

computing resources, AI-based classical approaches could simulate even the largest of these systems, says Tkatchenko. Given that quantum computers powerful enough to compete are likely still decades away, he adds, AI's current trajectory suggests it could reach important milestones, such as precisely simulating how drugs bind to a protein, much sooner.

Strong correlations

When it comes to simulating strongly correlated quantum systems—ones whose particles interact a lot—methods like DFT quickly run out of steam. While more exotic, these systems include materials with potentially transformative capabilities, like high-temperature superconductivity or ultra-precise sensing. But even here, AI is making significant strides.

In 2017, EPFL's Carleo and Microsoft's Troyer published a seminal [paper in *Science*](#) showing that neural networks could model strongly correlated quantum systems. The approach doesn't learn from data in the classical sense. Instead, Carleo says, it is similar to DeepMind's AlphaZero model, which mastered the games of Go, chess, and shogi using nothing more than the rules of each game and the ability to play itself.

In this case, the rules of the game are provided by Schrödinger's equation, which can precisely describe a system's quantum state, or wave function. The model plays against itself by arranging particles in a certain configuration and then measuring the system's energy level. The goal is to reach the lowest energy configuration (known as the ground state), which determines the system's properties. The model repeats this process until energy levels stop falling, indicating that the ground state—or something close to it—has been reached.

The power of these models is their ability to compress information, says Carleo. "The wave function is a very complicated mathematical object," he says. "What has been shown by several papers now is that [the neural network] is able to capture the complexity of this object in a way that can be handled by a classical machine."

Since the 2017 paper, the approach has been extended to a wide range of strongly correlated systems, says Carleo, and results have been impressive. The *Science* paper he published with colleagues last month put leading classical simulation techniques to the test on a variety of tricky quantum simulation problems, with the goal of creating a benchmark to judge advances in both classical and quantum approaches.



The industry's must-read newsletter ×


Understand what's happening in tech & why when you sign up for The Debrief with Mat Honan, our free newsletter.

SIGN UP

Carleo says that neural-network-based techniques are now the best approach for simulating many of the most complex quantum systems they tested. “Machine learning is really taking the lead in many of these problems,” he says.

These techniques are catching the eye of some big players in the tech industry. In August, researchers at DeepMind showed in a [paper in *Science*](#) that they could accurately model excited states in quantum systems, which could one day help predict the behavior of things like solar cells, sensors, and lasers. Scientists at Microsoft Research have also developed an [open-source software suite](#) to help more researchers use neural networks for simulation.

Related Story



Google says it's made a quantum computing breakthrough that reduces errors

The company's surface code technique allows its quantum bits to faithfully store and manipulate data for longer, which could pave the way for useful quantum computers.

One of the main advantages of the approach is that it piggybacks on massive investments in AI software and hardware, says [Filippo Vicentini](#), a professor of AI and condensed-matter physics at École Polytechnique in France, who was also a coauthor on the *Science* benchmarking paper: “Being able to leverage these kinds of technological advancements gives us a huge edge.”

There is a caveat: Because the ground states are effectively found through trial and error rather than explicit calculations, they are only approximations. But this is also why the approach could make progress on what has looked like an intractable problem, says Juan Carrasquilla, a researcher at ETH Zurich, and another coauthor on the *Science* benchmarking paper.

If you want to precisely track all the interactions in a strongly correlated system, the number of calculations you need to do rises exponentially with the system's size. But if you're happy with an answer that is just good enough, there's plenty of scope for taking shortcuts.

“Perhaps there's no hope to capture it exactly,” says Carrasquilla. “But there's hope to capture enough information that we capture all the aspects that physicists care about. And if we do that, it's basically indistinguishable from a true solution.”

And while strongly correlated systems are generally too hard to simulate classically, there are notable instances where this isn't the case. That includes some systems that are relevant for modeling high-temperature superconductors, according to a 2023 [paper in *Nature Communications*](#).

“Because of the exponential complexity, you can always find problems for which you can't find a shortcut,” says Frank Noe, research manager at Microsoft Research, who has led much of the company's work in this area. “But I think the number of systems for which you can't find a good shortcut will just become much smaller.”

No magic bullets

However, [Stefanie Czischek](#), an assistant professor of physics at the University of Ottawa, says it can be hard to predict what problems neural networks can feasibly solve. For some complex systems they do incredibly well, but then on other seemingly simple ones, computational costs balloon unexpectedly. “We don't really know their limitations,” she says. “No one really knows yet what are the conditions that make it hard to represent systems using these neural networks.”

Meanwhile, there have also been significant advances in other classical quantum simulation techniques, says [Antoine Georges](#), director of the Center for Computational Quantum Physics at the Flatiron Institute in New York, who also contributed to the recent *Science* benchmarking paper. “They are all successful in their own right, and they are also very complementary,” he says. “So I don't think these machine-learning methods are just going to completely put all the other methods out of business.”

THE CHECKUP ×

A weekly newsletter focusing on the most important news in health and biotech.

Enter your email

Sign up

By signing up, you agree to our [Privacy Policy](#).

Quantum computers will also have their niche, says [Martin Roetteler](#), senior director of quantum solutions at IonQ, which is developing quantum computers built from trapped ions. While he agrees that classical approaches will likely be sufficient for simulating weakly correlated systems, he's confident that some large, strongly correlated systems will be beyond their reach. “The exponential is going to bite you,” he says. “There are cases with strongly correlated systems that we cannot treat classically. I'm strongly convinced that that's the case.”

In contrast, he says, a future fault-tolerant quantum computer with many more qubits than today's devices will be able to simulate such systems. This could help find new catalysts or improve understanding of metabolic processes in the body—an area of interest to the

pharmaceutical industry.

Neural networks are likely to increase the scope of problems that can be solved, says [Jay Gambetta](#), who leads IBM's quantum computing efforts, but he's unconvinced they'll solve the hardest challenges businesses are interested in.

"That's why many different companies that essentially have chemistry as their requirement are still investigating quantum—because they know exactly where these approximation methods break down," he says.

Related Story



PsiQuantum plans to build the biggest quantum computing facility in the US

The company wants to build a computer containing up to 1 million qubits on a Chicago campus.

Gambetta also rejects the idea that the technologies are rivals. He says the future of computing is likely to involve a hybrid of the two approaches, with quantum and classical subroutines working together to solve problems. "I don't think they're in competition. I think they actually add to each other," he says.

But [Scott Aaronson](#), who directs the Quantum Information Center at the University of Texas, says machine-learning approaches are directly competing against quantum computers in areas like quantum chemistry and condensed-matter physics. He predicts that a combination of machine learning and quantum simulations will outperform purely classical approaches in many cases, but that won't become clear until larger, more reliable quantum computers are available.

"From the very beginning, I've treated quantum computing as first and foremost a scientific quest, with any industrial applications as icing on the cake," he says. "So if quantum simulation turns out to beat classical machine learning only rarely, I won't be quite as crestfallen as some of my colleagues."

One area where quantum computers look likely to have a clear advantage is in simulating how complex quantum systems evolve over time, says EPFL's Carleo. This could provide invaluable insights for scientists in fields like statistical mechanics and high-energy physics, but it seems unlikely to lead to practical uses in the near term. "These are more niche applications that, in my opinion, do not justify the massive investments and the massive hype," Carleo adds.

Nonetheless, the experts *MIT Technology Review* spoke to said a lack of commercial applications is not a reason to stop pursuing quantum computing, which could lead to fundamental scientific breakthroughs in the long run.

"Science is like a set of nested boxes—you solve one problem and you find five other problems," says Vicentini. "The complexity of the things we study will increase over time, so we will always need more powerful tools." **T**

by Edd Gent



NEW: The Food × issue

Join us as we take a closer look at the future of food and how we can use technology to feed more people.

SUBSCRIBE

TODAY

KEEP READING


MOST POPULAR



Meet the radio-obsessed civilian shaping Ukraine's drone defense

Since Russia's invasion, Serhii "Flash" Beskrestnov has become an influential, if sometimes controversial, force—sharing expert advice and intel on the ever-evolving technology that's taken over the skies. His work may determine the future of Ukraine, and wars far beyond it.

By Charlie Metcalfe



Why OpenAI's new model is such a big deal

The bulk of LLM progress until now has been language-driven. This new model enters the realm of complex reasoning, with implications for physics, coding, and more.

By James O'Donnell



Two Nobel Prize winners want to cancel their own CRISPR patents in Europe

There's a surprise twist in the battle to control genome editing.

By Antonio Regalado



Why Microsoft made a deal to help restart Three Mile Island

A once-shuttered nuclear plant could soon return to the grid.

By Casey Crownhart

STAY CONNECTED

Illustration by Rose Wong

Get the latest updates from MIT Technology Review

Discover special offers, top stories, upcoming events, and more.

Enter your email

[Privacy Policy](#)

The latest iteration of a legacy

Founded at the Massachusetts Institute of Technology in 1899, MIT Technology Review is a world-renowned, independent media company whose insight, analysis, reviews, interviews and live events explain the newest technologies and their commercial, social and political impact.

Advertise with MIT Technology Review

Elevate your brand to the forefront of conversation around emerging technologies that are radically transforming business. From event sponsorships to custom content to visually arresting video storytelling, advertising with MIT Technology Review creates opportunities for your brand to resonate with an unmatched audience of technology and business elite.

READ ABOUT OUR HISTORY

ADVERTISE WITH US

[About us](#)

[Help & FAQ](#)

[Careers](#)

[My subscription](#)

[Custom content](#)

[Editorial guidelines](#)

[Advertise with us](#)

[Privacy policy](#)

[International Editions](#)

[Terms of Service](#)

[Republishing](#)

[Write for us](#)

[MIT Alumni News](#)

[Contact us](#)

