


How Artificial Intelligence Is Changing Science

The latest AI algorithms are probing the evolution of galaxies, calculating quantum wave functions, discovering new chemical compounds and more. Is there anything that scientists do that can't be automated?

[Spoiler alert: answer on last page.]

Dan Falk, 11 March 2019, *Quanta Magazine*
<https://www.quantamagazine.org/how-artificial-intelligence-is-changing-science-20190311>



No human, or team of humans, could possibly keep up with the avalanche of information produced by many of today's physics and astronomy experiments. Some of them record terabytes of data every day — and the torrent is only increasing. The Square Kilometer Array, a radio telescope slated to switch on in the mid-2020s, will generate about as much data traffic each year as the entire internet.

The deluge has many scientists turning to artificial intelligence for help. With minimal human input, AI systems such as artificial neural networks — computer-simulated networks of neurons that mimic the function of brains — can plow through mountains of data, highlighting anomalies and detecting patterns that humans could never have spotted.

Of course, the use of computers to aid in scientific research goes back about 75 years, and the method of manually poring over data in search of meaningful patterns originated

millennia earlier. But some scientists are arguing that the latest techniques in machine learning and AI represent a fundamentally new way of doing science. One such approach, known as generative modeling, can help identify the most plausible theory among competing explanations for observational data, based solely on the data, and, importantly, without any preprogrammed knowledge of what physical processes might be at work in the system under study. Proponents of generative modeling see it as novel enough to be considered a potential “third way” of learning about the universe.

Traditionally, we've learned about nature through observation. Think of Johannes Kepler poring over Tycho Brahe's tables of planetary positions and trying to discern the underlying pattern. (He eventually deduced that planets move in elliptical orbits.) Science has also advanced through simulation. An astronomer might model the movement of the Milky Way and its neighboring galaxy, Andromeda, and predict that they'll collide in a few billion years. Both observation and simulation help scientists generate hypotheses that can then be tested with further observations. Generative modeling differs from both of these approaches.

“It's basically a third approach, between observation and simulation,” says Kevin Schawinski, an astrophysicist and one of generative modeling's most enthusiastic proponents, who worked until recently at the Swiss Federal Institute of Technology in Zurich (ETH Zurich). “It's a different way to attack a problem.”

Some scientists see generative modeling and other new techniques simply as power tools for doing traditional science. But most agree that AI is having an enormous impact, and that its role in science will only grow. Brian Nord, an astrophysicist at Fermi National Accelerator Laboratory who uses artificial neural networks to study the cosmos, is among those who fear there's nothing a human scientist does that will be impossible to automate. “It's a bit of a chilling thought,” he said.

Discovery by Generation

Ever since graduate school, Schawinski has been making a name for himself in data-driven science. While working on his doctorate, he faced the task of classifying thousands of galaxies based on their appearance. Because no readily available software existed for the job, he decided to crowdsource it — and so the Galaxy Zoo citizen science project was born. Beginning in 2007, ordinary computer users helped astronomers by logging their best guesses as to which galaxy belonged in which category, with majority rule typically leading to correct classifications. The project was a success, but, as Schawinski notes, AI has made it obsolete: “Today, a talented scientist with a background in machine learning and access to cloud computing could do the whole thing in an afternoon.”

Schawinski turned to the powerful new tool of generative modeling in 2016. Essentially, generative modeling asks how likely it is, given condition X, that you’ll observe outcome Y. The approach has proved incredibly potent and versatile. As an example, suppose you feed a generative model a set of images of human faces, with each face labeled with the person’s age. As the computer program combs through these “training data,” it begins to draw a connection between older faces and an increased likelihood of wrinkles. Eventually it can “age” any face that it’s given — that is, it can predict what physical changes a given face of any age is likely to undergo.



None of these faces is real. The faces in the top row (A) and left-hand column (B) were constructed by a generative adversarial network (GAN) using building-block elements of real faces. The GAN then combined basic features of the faces in A, including their gender, age and face shape, with finer features of faces in B, such as hair color and eye color, to create all the faces in the rest of the grid.

Adapted from NVIDIA

The best-known generative modeling systems are “generative adversarial networks” (GANs). After adequate exposure to training data, a GAN can repair images that have damaged or missing pixels, or they can make blurry photographs sharp. They learn to infer the missing information by means of a competition (hence the term “adversarial”): One part of the network, known as the generator, generates fake data, while a second part, the discriminator, tries to distinguish fake data from real data. As the program runs, both halves get progressively better. You may have seen some of the hyper-realistic, GAN-produced “faces” that have circulated recently — images of “freakishly realistic people who don’t actually exist,” as one headline put it.

More broadly, generative modeling takes sets of data (typically images, but not always) and breaks each of them down into a set of basic, abstract building blocks — scientists refer to this as the data’s “latent space.” The algorithm manipulates elements of the latent space to see how this affects the original data, and this helps uncover physical processes that are at work in the system.

The idea of a latent space is abstract and hard to visualize, but as a rough analogy, think of what your brain might be doing when you try to determine the gender of a human face. Perhaps you notice hairstyle, nose shape, and so on, as well as patterns you can’t easily put into words. The computer program is similarly looking for salient features among data: Though it has no idea what a mustache is or what gender is, if it’s been trained on data sets in which some images are tagged “man” or “woman,” and in which some have a “mustache” tag, it will quickly deduce a connection.


In a paper published in December in *Astronomy & Astrophysics*, Schawinski and his ETH Zurich colleagues Dennis Turp and Ce Zhang used generative modeling to investigate the physical changes that galaxies undergo as they evolve. (The software they used treats the latent space somewhat differently from the way a generative adversarial network treats it, so it is not technically a GAN, though similar.) Their model created artificial data sets as a way of testing hypotheses about physical processes. They asked, for instance, how the “quenching” of star formation — a sharp reduction in formation rates — is related to the increasing density of a galaxy’s environment.

For Schawinski, the key question is how much information about stellar and galactic processes could be teased out of the data alone. “Let’s erase everything we know about astrophysics,” he said. “To what degree could we rediscover that knowledge, just using the data itself?”

First, the galaxy images were reduced to their latent space; then, Schawinski could tweak one element of that space in a way that corresponded

to a particular change in the galaxy’s environment — the density of its surroundings, for example. Then he could re-generate the galaxy and see what differences turned up. “So now I have a hypothesis-generation machine,” he explained. “I can take a whole bunch of galaxies that are originally in a low-density environment and make them look like they’re in a high-density environment, by this process.” Schawinski, Turp and Zhang saw that, as galaxies go from low- to high-density environments, they become redder in color, and their stars become more centrally concentrated. This matches existing observations about galaxies, Schawinski said. The question is why this is so.

The next step, Schawinski says, has not yet been automated: “I have to come in as a human, and say, ‘OK, what kind of physics could explain this effect?’” For the process in question, there are two plausible explanations: Perhaps galaxies become redder in high-density environments because they contain more dust, or perhaps they become redder because of a decline in star formation (in other words, their stars tend to be older). With a generative model, both ideas can be put to the test: Elements in the latent space related to dustiness and star formation rates are changed to see how this affects galaxies’ color. “And the answer is clear,” Schawinski said. Redder galaxies are “where the star formation had dropped, not the ones where the dust changed. So we should favor that explanation.”



Using generative modeling, astrophysicists could investigate how galaxies change when they go from low-density regions of the cosmos to high-density regions, and what physical processes are responsible for these changes.

Adapted from K. Schawinski et al.; Source doi: 10.1051/0004-6361/201833800

The approach is related to traditional simulation, but with critical differences. A simulation is “essentially assumption-driven,” Schawinski said. “The approach is to say, ‘I think I know what the underlying physical laws are that give rise to everything that I see in the system.’ So I have a recipe for star formation, I have a recipe for how dark matter behaves, and so on. I put all of my hypotheses in there, and I let the simulation run. And then I ask: Does that look like reality?” What he’s done with generative modeling, he said, is “in some sense, exactly the opposite of a simulation. We don’t know anything; we don’t want to assume anything. We want the data itself to tell us what might be going on.”

The apparent success of generative modeling in a study like this obviously doesn’t mean that astronomers and graduate students have been made redundant — but it appears to represent a shift in the degree to which learning about astrophysical objects and processes can be achieved by an artificial system that has little more at its electronic fingertips than a vast pool of data. “It’s not fully automated science — but it demonstrates that we’re capable of at least in part building the tools that make the process of science automatic,” Schawinski said.

Generative modeling is clearly powerful, but whether it truly represents a new approach to science is open to debate. For David Hogg, a cosmologist at New York University and the Flatiron Institute (which, like *Quanta*, is funded by the Simons Foundation), the technique is impressive but ultimately just a very sophisticated way of extracting patterns from data — which is what astronomers have been doing for centuries. In other words, it’s an advanced form of observation plus analysis. Hogg’s own work, like Schawinski’s, leans heavily on AI; he’s been using neural networks to classify stars according to their spectra and to infer other physical attributes of stars using data-driven models. But he sees his work, as well as Schawinski’s, as tried-and-true science. “I don’t think it’s a third way,” he said recently. “I just think we as a community are becoming far more sophisticated about how we use the data. In particular, we are getting much better at comparing data to data. But in my view, my work is still squarely in the observational mode.”

Hardworking Assistants

Whether they’re conceptually novel or not, it’s clear that AI and neural networks have come to play a critical role in contemporary astronomy and physics research. At the Heidelberg Institute for Theoretical Studies, the physicist Kai Polsterer heads the astroinformatics group — a team of researchers focused on new, data-centered methods of doing astrophysics. Recently, they’ve been using a machine-learning algorithm to extract redshift information from galaxy data sets, a previously arduous task.

Polsterer sees these new AI-based systems as “hardworking assistants” that can comb through data for hours on end without getting bored or complaining about the working conditions. These systems can do all the tedious grunt work, he said, leaving you “to do the cool, interesting science on your own.”

But they’re not perfect. In particular, Polsterer cautions, the algorithms can only do what they’ve been trained to do. The system is “agnostic” regarding the input. Give it a galaxy, and the software can estimate its redshift and its age — but feed that same system a selfie, or a picture of a rotting fish, and it will output a (very wrong) age for that, too. In the end, oversight by a human scientist remains essential, he said. “It comes back to you, the researcher. You’re the one in charge of doing the interpretation.”

For his part, Nord, at Fermilab, cautions that it’s crucial that neural networks deliver not only results, but also error bars to go along with them, as every undergraduate is trained to do. In science, if you make a measurement and don’t report an estimate of the associated error, no one will take the results seriously, he said.

Like many AI researchers, Nord is also concerned about the impenetrability of results produced by neural networks; often, a system delivers an answer without offering a clear picture of how that result was obtained.

Yet not everyone feels that a lack of transparency is necessarily a problem. Lenka Zdeborová, a researcher at the Institute of Theoretical Physics at CEA Saclay in France, points out that human intuitions are often equally impenetrable. You look at a photograph and instantly recognize a cat — “but you don’t know how you know,” she said. “Your own brain is in some sense a black box.”

It’s not only astrophysicists and cosmologists who are migrating toward AI-fueled, data-driven science. Quantum physicists like Roger Melko of the Perimeter Institute for Theoretical Physics and the University of Waterloo in Ontario have used neural networks to solve some of the

toughest and most important problems in that field, such as how to represent the mathematical “wave function” describing a many-particle system. AI is essential because of what Melko calls “the exponential curse of dimensionality.” That is, the possibilities for the form of a wave function grow exponentially with the number of particles in the system it describes. The difficulty is similar to trying to work out the best move in a game like chess or Go: You try to peer ahead to the next move, imagining what your opponent will play, and then choose the best response, but with each move, the number of possibilities proliferates.

Of course, AI systems have mastered both of these games — chess, decades ago, and Go in 2016, when an AI system called AlphaGo defeated a top human player. They are similarly suited to problems in quantum physics, Melko says.

The Mind of the Machine

Whether Schawinski is right in claiming that he’s found a “third way” of doing science, or whether, as Hogg says, it’s merely traditional observation and data analysis “on steroids,” it’s clear AI is changing the flavor of scientific discovery, and it’s certainly accelerating it. How far will the AI revolution go in science?

Occasionally, grand claims are made regarding the achievements of a “robo-scientist.” A decade ago, an AI robot chemist named Adam investigated the genome of baker’s yeast and worked out which genes are responsible for making certain amino acids. (Adam did this by observing strains of yeast that had certain genes missing, and comparing the results to the behavior of strains that had the genes.) *Wired*’s headline read, “Robot Makes Scientific Discovery All by Itself.”

More recently, Lee Cronin, a chemist at the University of Glasgow, has been using a robot to randomly mix chemicals, to see what sorts of new compounds are formed. Monitoring the reactions in real-time with a mass spectrometer, a nuclear magnetic resonance machine, and an

infrared spectrometer, the system eventually learned to predict which combinations would be the most reactive. Even if it doesn't lead to further discoveries, Cronin has said, the robotic system could allow chemists to speed up their research by about 90 percent.

Last year, another team of scientists at ETH Zurich used neural networks to deduce physical laws from sets of data. Their system, a sort of robo-Kepler, rediscovered the heliocentric model of the solar system from records of the position of the sun and Mars in the sky, as seen from Earth, and figured out the law of conservation of momentum by observing colliding balls. Since physical laws can often be expressed in more than one way, the researchers wonder if the system might offer new ways — perhaps simpler ways — of thinking about known laws.

These are all examples of AI kick-starting the process of scientific discovery, though in every case, we can debate just how revolutionary the new approach is. Perhaps most controversial is the question of how much information can be gleaned from data alone — a pressing question in the age of stupendously large (and growing) piles of it. In *The Book of Why* (2018), the computer scientist Judea Pearl and the science writer Dana Mackenzie assert that data are “profoundly dumb.” Questions about causality “can never be answered from data alone,” they write. “Anytime you see a paper or a study that analyzes the data in a model-free way, you can be certain that the output of the study will merely summarize, and perhaps transform, but not interpret the data.” Schawinski sympathizes with Pearl's position, but he described the idea of working with “data alone” as “a bit of a straw man.” He's never claimed to deduce cause and effect that way, he said. “I'm merely saying we can do more with data than we often conventionally do.”

Another oft-heard argument is that science requires creativity, and that — at least so far — we have no idea how to program that into a machine. (Simply trying everything, like Cronin's robo-chemist, doesn't seem especially creative.) “Coming up with a theory, with reasoning, I think demands creativity,” Polsterer said. “Every time you need creativity, you will need a human.” And where does creativity come from? Polsterer suspects it is related to boredom — something that, he says, a machine cannot experience. “To be creative, you have to dislike being bored. And I don't think a computer will ever feel bored.” On the other hand, words like “creative” and “inspired” have often been used to describe programs like Deep Blue and AlphaGo. And the struggle to describe what goes on inside the “mind” of a machine is mirrored by the difficulty we have in probing our own thought processes.

Schawinski recently left academia for the private sector; he now runs a startup called Modulos which employs a number of ETH scientists and, according to its website, works “in the eye of the storm of developments in AI and machine learning.” Whatever obstacles may lie between current AI technology and full-fledged artificial minds, he and other experts feel that machines are poised to do more and more of the work of human scientists. Whether there is a limit remains to be seen.

“Will it be possible, in the foreseeable future, to build a machine that can discover physics or mathematics that the brightest humans alive are not able to do on their own, using biological hardware?” Schawinski wonders. “Will the future of science eventually necessarily be driven by machines that operate on a level that we can never reach? I don't know. It's a good question.”