



ANTONELLO VENERI/AFP VIA GETTY

The impact of climate change on Brazil's Atlantic coastline is a research focus at the University of São Paulo's machine-intelligence centre.

# AI AND ITS GROWING FOOTHOLD IN THE NATURAL SCIENCES

Collaborations across disciplines are growing, with lessons learnt on how to make joint working more effective.

**A**rtificial intelligence (AI) is increasingly becoming a tool for researchers in other science and technology fields, forging collaborations across disciplines. Stanford University in California, which produces an index that tracks AI-related data, finds in its 2021 report that the number of AI journal publications grew by 34.5% from 2019 to 2020; up from 19.6% between 2018 and 2019 (see [go.nature.com/3mdt2yq](https://go.nature.com/3mdt2yq)). AI publications represented 3.8% of all peer-reviewed scientific publications worldwide in 2019, up from 1.3% in 2011.

Five AI researchers describe the fruits of these collaborations, beyond journal publications, and talk about how they are helping to break down barriers between disciplines.

## FABIO COZMAN MANAGE EXPECTATIONS

At the University of São Paulo in Brazil, where I lead the Center for Artificial Intelligence (C4AI), our main goal is to produce machine-intelligence research that has a direct impact on society and industry. We have five core programmes. One aims to greatly improve natural-language processing and translation in Portuguese, the language of Brazil, so that what Portuguese speakers say can be translated, transcribed and understood much better by computerized speech tools. Another, the Blue Amazonia Brain, examines the influence of climate change, biodiversity and mineral

resources on Brazil's Atlantic coastline and the people who live there. The centre opened in October 2020, with annual funding of 2 million Brazilian reais (\$US380,000) from the technology company IBM, 2 million reais from the São Paulo Research Foundation and 4 million reais from the University of São Paulo. The state government provides further financial support.

The centre collaborates widely, but collaborators often have different expectations about what computer science can achieve. These expectations can be addressed by being clear with collaborators about what AI can and can't (yet) do. And disagreements frequently arise about research outputs: for example, people in the natural sciences generally see

journal papers as the best way of disseminating research, whereas, in my experience, AI researchers put more value in conferences.

Another challenge is that some researchers just want a programmer. Such researchers need to be more willing to share their knowledge and problems, rather than just adopting a ‘come and help me do programming’ approach. Collaboration needs to be a partnership that aims to address and answer questions.

AI has grown so fast that people in computer science and engineering feel they have to reach out to solve real-world problems: just doing our own thing is no longer that rewarding for us. We’re following a trend: all major AI laboratories and centres are now getting involved in real-life, applied problems. My advice for researchers who are hoping to collaborate with AI specialists is to first manage your expectations: Are you hoping to have someone who’s good at computers’ help you to do some data analysis, or do you actually need to ask much deeper questions, which AI might be able to help you answer?

A little background knowledge and practical experience is useful among collaborators.

**Fabio Cozman** is director of the Center for Artificial Intelligence (C4AI) at the University of São Paulo, Brazil.

## PHIALA SHANAHAN OPERATE ON AN EVEN FOOTING

I have an ongoing collaboration with Google DeepMind, the company’s AI research division. This association started at a conference in Israel a couple of years ago. My students and I presented a few projects that we had started at the Massachusetts Institute of Technology (MIT) in Cambridge, using some ideas developed by Danilo Jimenez Rezende, a senior staff research scientist at DeepMind in London; Rezende’s work includes the modelling of complex data such as medical images, videos, 3D-scene geometry and complex physical systems. He had done some key machine-learning research that we had already applied to problems in fundamental physics.

We spoke, and a longer-term collaboration has grown out of that. It now involves several people at DeepMind, a couple of my postdocs and a PhD student. We’ve written four or five papers over the past few years and have really done some innovative things, using machine-learning models to accelerate established physics calculations. Ultimately, the goal is to enable us to undertake studies that are computationally impossible with existing algorithms and resources.



**Theoretical physicist Phiala Shanahan.**

Something that makes our collaboration successful is the sense of equality. My group is pushing on the AI side as hard as the DeepMind group is. And the people in the DeepMind group seriously know their physics, too. Both sides can do both parts of the science, so it has been a really even and dynamic collaboration, and really good fun.

I’ve been involved in less effective collaborations, where the attitude is that ‘one group should worry about the physics part, and one group should worry about the computer science part’, and we meet in the middle. What happens is that both groups end up getting siloed and fighting a language barrier. I’ve found such cooperations not to be interactive.

In practice, what this closer, more even relationship with DeepMind means is that we have a meeting once a week with everyone involved in the project. We also have a joint channel on the collaboration platform Slack, where we chat in the meantime, and I have meetings more frequently during the week with those in my own group who are working on the project.

**Phiala Shanahan** is a theoretical physicist at the Massachusetts Institute of Technology in Cambridge.

## SIMON OLSSON FIND A PROBLEM TO SOLVE

In my lab, which I started in October last year, we develop machine-learning methods to solve computational problems in the natural sciences. At the moment, for example, we’re developing methods for designing pharmaceuticals in collaboration with the British–Swedish company AstraZeneca, which has a research centre in Gothenburg, near my

university. We’re also working on ways to integrate experimental data into machine-learning models of protein structure and dynamics.

We use published papers and data from the natural sciences to train algorithms, rather than letting them work things out on their own. For example, if you’re trying to work out how a protein folds, or how a drug interacts with it, then using a computational model that takes into account the literature about that protein, as well as the physical and chemical laws that govern how it behaves, will probably be helpful.

If you have a computer-science background and want to get into AI while studying natural sciences, try to identify an area that you’re interested in and find a problem that you’d like to solve. For example, I first became drawn to this field through studying molecular dynamics and molecular design, in which molecules and their interactions are simulated in a computer, often for drug-discovery purposes. AI has the potential to make previously unsolvable problems solvable in fields such as these, which are enormously demanding computationally.

If you don’t come from a computer-science background, it’s important to learn to program and to get to grips with the fundamentals of machine-learning theory. One place to start is [learnpython.org](http://learnpython.org), which provides an interactive tutorial on the programming language Python. There are also online courses on machine learning on the US online course platform Coursera, and on YouTube. Or you could attend a course on machine learning or data science at your university.

Picking up the basics of programming with AI also means developing knowledge of applied statistics and studying how machine-learning algorithms work, as well as some of the ways in which they process data and ‘learn’ from experience. Getting a grasp of those concepts is an important first step.

I think that recognizing the usefulness of machine learning and AI really comes down to asking yourself: ‘How can these methods help us improve, to push science forward in a fundamental way?’

I’d advise someone who’s interested in AI to start learning programming by simply trying to automate something that they do regularly in their working lives: whether that’s sending a templated e-mail or entering data into a spreadsheet. If it’s a boring task to repeat, then the motivation to automate it will come really quickly. After that, gradually challenge yourself with more and more complicated tasks.

**Simon Olsson** is an assistant professor of applied artificial intelligence at Chalmers University of Technology in Gothenburg, Sweden.

## ROMAN LIPSKI MY MACHINE-LEARNING MUSE

In April 2016, I started teaching a course to refugees at the College of Fine Arts in Berlin, where I met Florian Dohmann, a data scientist. We started a collaboration to try to explore art using artificial intelligence. I'd seen pictures that data scientists at Google had made using AI. They were horror-story images made from repeating elements, featuring animals with 1,000 eyes or 1,000 feet.

At first I was a little naive: I thought perhaps we'd immediately develop the best picture ever. I knew there was huge potential, but I didn't know how to use AI. Florian and I started by working with an open-source algorithm, created by scientists at the University of Tübingen in Germany, which was designed to recognize shapes and colours using machine learning. To stay true to my artistic principles, we decided to train the algorithm only on my own work. We photographed every painting I'd made in my career, to create a small data set with which to teach the algorithm, and then asked it to create an original piece of work.

The result was again horrible. It looked like the paintings I had seen from the Google engineers – repeating shapes and colours without adding anything new. Artistically speaking, they were more gimmick than anything else.

We decided to create a new, data set, using a repeating motif – inspired by Andy Warhol's *Campbell's Soup Cans* – that I'd used in my own



Artist Roman Lipski uses AI as inspiration.

paintings. The motif was a very simple landscape of a street in Los Angeles that I'd visited in March 2016. I'd painted that same scene several times, in different colours and textures.

We digitized that set of images, and I realized that I was making art not to be exhibited to a human audience, but rather to be 'seen' and processed by a machine: it was the start of a dialogue between me and the machine.

This time, when we asked the algorithm to innovate and make new pictures, the results were amazing. Not every picture was good, but we got thousands and thousands of great results in different artistic styles, with real artistic quality and in forms that I would not have arrived at by myself.

A year before meeting Florian, I had hit a complete artistic crisis; I felt I'd run out of stories from my own world that I could tell in paint. I've now started painting again, but rather than simply printing what the AI algorithm generates, I use its output as inspiration to create my own original works. I now encourage others to use the algorithm as part of a community art project called Unfinished, helping them to experience my creative process with the AI tool and create their own paintings.


My advice is to not be intimidated by AI devices but just to start using them: like any tools, they have their upsides and their downsides. But for me, AI changed my career for the better.

**Roman Lipski** is a Berlin-based artist who incorporates AI into his work.

## SIDDHARTH MISHRA-SHARMA FIND GREAT MENTORS

I did some internships in experimental high-energy physics and astrophysics as an undergraduate at the University of Cambridge, UK. They included a couple of summers at CERN, Europe's particle-physics laboratory near Geneva, Switzerland. I also dabbled with machine learning during my PhD in particle physics at Princeton University, New Jersey, and have returned to it in my current role at MIT. AI tools tend to be a great complement to physics. We often work with huge data sets from particle colliders or telescopes, which can produce petabytes of data.

For example, suppose you have a vast data set from tracking the movements of stars through our Galaxy. Dark matter can have interesting effects on the motions of stars, pulling them one way or the other slightly, or distorting the light that comes from them. Because the effect is subtle, it's hard to analyse 100 billion-plus stars individually. Ultimately, it becomes a big-data



Particle physicist Siddharth Mishra-Sharma.

problem: machine-learning methods can help us to recognize patterns, and can be scaled to handle huge data sets.

And because so many types of astronomical data set are available – from images of individual galaxies to maps of the Milky Way – no single machine-learning method can be used effectively to look for the effects of dark matter. When machine learning began to be used in astrophysics, methods were adapted wholesale, using established algorithms in new contexts. If a machine-learning method was good at distinguishing between images of cats and dogs, for example, it would be adapted to distinguish between images of different galaxies.

But today, the needs of physicists and other practitioners in the natural sciences can inform the development of machine-learning methods. I no longer work directly with massive data sets straight from colliders or telescopes. Instead, part of my day-to-day work involves looking at what kinds of method would work well for a given problem or observation and, if no such method exists, trying to create it. In work such as this, the flow of information between physics and machine learning is moving in both directions, and the two disciplines are informing each other. I'm excited to be part of this.

I'd encourage others to reach out to potential mentors and say: "Here's an interesting problem – I think your method is perfect for it." Often, those on the other side of a collaboration are excited about modifying their method to suit your needs, or providing advice. They're often only too happy to think about your problem.

**Siddharth Mishra-Sharma** is a postdoctoral particle-physics researcher at the Massachusetts Institute of Technology in Cambridge.

### Interviews by Jack Leeming.

These interviews have been edited for length and clarity.