

Bringing Quantum to Machine Learning

Maria Schuld reflects on the open questions about quantum machine-learning algorithms.

By **Katherine Wright**

Be-coming a physicist was not Maria Schuld's life goal. As an undergrad, she started out studying political science, taking physics in parallel. Her plan was to work for a nonprofit organization in a capacity that had a very clear benefit to society. For her, a career as a physicist didn't offer that possibility. But then, she says, "life happened"—jobs fell through and other opportunities opened up—and she found herself with a career in quantum machine learning.



Credit: Z. Giampietri

Today Schuld, who works for the Canadian quantum computing company Xanadu from her home in South Africa, says that she has matured in what she thinks it means for a person to benefit society. She says that all people can take actions to benefit others, regardless of their field. For example, any startup company can build a supportive culture with happy employees or a ruthless one that makes everyone miserable. In an interview with *Physics*, Schuld spoke about why she loves quantum machine learning, what she sees as the important unsolved problems in the field, and how she approaches career decisions.

All interviews are edited for brevity and clarity.

What is a quantum machine-learning model?

That's a question with no one answer. For me, a quantum machine-learning model is one where the thing that's used to solve the task is a quantum computation. These computations don't have clear recipes to follow, like Shor's algorithm—a quantum algorithm for integer factorization. Rather, they are more an abstract skeleton that the model uses to train itself.

What quantum machine-learning models are you interested in?

I absolutely love models that use so-called kernel methods. Classical kernel methods are a class of algorithm used for pattern analysis, and they were very popular in the 1990s. The mathematics of quantum computing looks very similar to classical kernel methods. This similarity allows us to apply results from classical methods to quantum computing. I find this similarity really interesting. Mathematically, something cool is happening.

Why use a quantum machine-learning algorithm over a classical one?

That's another question with many answers. Speedup is one goal—you probably wouldn't use a quantum algorithm if it wasn't faster. But I'm more interested in whether a model will perform better if we replace it with a generic quantum computation. How to answer that question is much less clear because we don't have a good sense of how to define "better."

The holy grail of machine learning is generalization power—the ability to apply the same model to different situations. For example, you might want to train a model to play a specific game and then use the same model to play a completely different game. Knowing if a quantum machine-learning algorithm generalizes is a really hard problem, as we don't have the theoretical tools we need to solve that problem.

Classical models do already generalize pretty well, right?

Yes. That is what machine-learning models, such as neural networks, do super well. Today you can train a neural network on a million images, and then give it a million unseen images, and the model will correctly tell you what is in every image.

What we still don't know about classical models is the ideal size. Initially, researchers thought that a model that was neither super small nor super big in terms of the number of parameters would be the best choice to optimize learning. And there are lots of theories that explain why that should be the case. But then they tried making the models super big and found that the learning ability just got better and better.

Do quantum machine-learning models show the same

improvement?

We don't know yet. We don't yet have the hardware and simulation abilities to test that question, and the theory on that problem is very thin. That is why everyone has been more focused on studying speedup—that is something we can do.

What is next for the field?

For me, it's better theory. We are very far from being able to do meaningful experiments. To recognize a typical image today would require millions of quantum gates, yet the best experiments have just a handful of gates. With theory, we can build models to answer how quantum machine-learning algorithms might work and what improvements they might show. Then, when the machines are ready, we can start testing the answers.

What advice would you give someone starting out in their career either in physics or in any other field?

Make career decisions based on what makes you happy. Early on, I was listening too much to other people's opinions when making big life decisions. And I can see that many of the young scientists that I mentor do the same. They make decisions based on what others think they should do, and these decisions can be bad for their mental health and bad for their futures. I turned down jobs in the US and in Europe that others would consider prestigious, because they didn't allow me to work from South Africa. I would give up physics before I give up living in a place I love.

Katherine Wright is a Senior Editor for *Physics*.