Quantum machine learning is going to be the biggest application of quantum computing in the next ten years

Prof. Peter Wittek is a thought leader in quantum-enhanced machine learning, quantum many-body systems, optimization, and machine learning, six years after he ended-up working by chance with quantum physicists. He wrote a book Quantum Machine Learning: What Quantum Computing Means to Data Mining.

Peter Wittek has been trained as a mathematician and a computer scientist. He received his PhD from the National University of Singapore.

He holds senior positions in several academic and research programs:

- Assistant Professor in the University of Toronto working on quantum-enhanced machine learning and applications of high-performance learning algorithms in quantum physics.
- Academic Director of the Quantum Machine Learning Program in the Creative Destruction Lab a seed-stage program for massively scalable, science-based companies.
- Faculty Affiliate in the Vector Institute for Artificial Intelligence a leader in the transformative field of artificial intelligence, excelling in machine and deep learning.
- Affiliate in the Perimeter Institute for Theoretical Physics a leading center for scientific research, training and educational outreach in foundational theoretical physics.

In this interview by Alain Chancé, Peter Wittek details his vision of quantum-enhanced machine learning in the next ten years.

1/ What is your vision of quantum-enhanced machine learning in the next 10 years?

Until scalable fault-tolerant universal quantum computers are available in maybe 20 years, quantum computing will be based on noisy quantum processors. Fortunately, machine learning also likes noise. In the next ten years, quantum machine learning is going to be the biggest application of quantum computing.

We already use a range of GPU's (Graphics Processing Units) to run deep learning networks. In the last few years, application-specific integrated circuits such as neuromorphic chips and tensor processing units have been added to the mix of AI hardware. Quantum computers will be part of this heterogeneous mix. In fact, I believe that for advancing AI, we must add quantum hardware to the mix of architecture.

Quantum computers can give a lift in probabilistic graphical models, Bayesian networks and Markov random fields, replacing Markov chains in Monte Carlo methods, and thus they could also boost Boltzmann and Helmholtz machines. Furthermore, they excel at discrete optimization, which is important in ensemble learning and metalearning, that is learning to learn. These are the most basic applications today for solving non-native quantum problems, that is, when you start with classical data and you want classical data in the end, with some quantum processing in between.

2/ What are the most important priorities right now?

The priority is educating the public, telling people that exponential speedup is not going to happen anytime soon, educating high school kids and training a whole new set of quantum computing developers. We badly need quantum computing experts. They are in very short supply.

The other priority is making advances with the hardware. Currently, there are four paradigms and all of them are noisy quantum processors that could all see significant advances:

1. Gate model quantum computing, currently the largest chip has around 70 qubits, there are a number of research programs: IBM, Google, Intel, Microsoft, Rigetti, SQC, and IonQ. This paradigm is the generalization of digital computing: instead of logical gates operating on bits, we have unitary gates operating on qubits, transforming quantum states to quantum states.
2. Quantum annealing. The D-Wave 2000Q™ quantum computer has 2,048 qubits. This paradigm uses a phenomenon in quantum physics to solve optimization problems, but it also works well for sampling certain families of probability distributions.
3. Gate model quantum computing using continuous variables, which uses qumodes instead of qubits. Xanadu designs and integrates quantum silicon photonic chips which

features qumodes. Unlike a qubit, which is a superposition of zero and one, a qumode is a superposition of real values.

4. Quantum simulation, one of the largest yet built is Harvard-MIT team's 51-Qubit Quantum Simulator. It is programmed by capturing super-cooled rubidium atoms with lasers and arranging them in a specific order, then allowing quantum mechanics to do the necessary calculations. Quantum simulators were the original motivation behind quantum computing, and we can think of them as the application-specific integrated circuits of the quantum world.

In digital computing there are currently few providers. Intel dominates in PCs-CPUs with AMD a distant second and only competitor. I believe that the third priority is to avoid creating monopolies in quantum computing. Now, in North America there are several companies in two countries in the quantum computing market. The European Union is going to wake-up. Dutch researchers collaborate with Intel. Australia and France have a joint venture between Silicon Quantum Computing in Sydney (SQC) and France's Commissariat à l'Energie Atomique et aux Energies Alternatives (CEA) to develop the world's first commercial silicon quantum integrated circuit.

3/ What are practical applications of quantum-enhanced machine learning?

A typical quantum native application is quantum chemistry computation. For the Finance Industry, the two most important applications are quantum-enhanced sampling and discrete optimization.

<u>Quantum-enhanced sampling</u>
Measuring a quantum system is essentially pulling-out a sample of a probability distribution. Early quantum devices could replace Markov chains in Monte Carlo methods with applications in asset pricing and option pricing. Several Canadian financial institutions are excited about exploring this domain.

<u>Discrete optimization</u>
In Finance, given a bunch of financial assets, how do you optimize the yield of this asset? This is an optimization problem. Shallow learning methods are typically not accurate. Let us take an example far detached from the finance industry: initially, deep learning methods were used to discover the Higgs boson, a massive computational and pattern recognition challenge. Last year, researchers managed to get a shallow learning method to rediscover the Higgs boson using a quantum annealing for machine learning approach that enabled them to extract physical insight [1]. If you translate this back to the finance industry, you could end up with more explainable models in a variety of applications.

1. A. Mott, J. Job, J. R. Vlimant, D. A. Lidar, and M. Spiropulu, “Solving a Higgs optimization problem with quantum annealing for machine learning”, Nature 550, 375 (2017), https://doi.org/10.1038/nature24047

4/ With whom are you collaborating?

I collaborate with a wide range of institutions, both in academia and in industry. My academic network spans six continents, including Africa, Asia, Latin America.

I also collaborate with many different hardware vendors, primarily with D-Wave, Rigetti, and Xanadu, and the team I am in has contacts to several others.

As the academic director of the world’s first QML startup incubator, I see about 25 companies going through the program every year, among them Adaptive Finance Technologies, based in Toronto, which pioneers Quantum Machine Learning technologies in the areas of financial services and investment management.

We also keep reaching out to corporations interested in quantum computing, such as banks, automotive and pharma companies.