The Quantum Leap
February 22, 2022

At the Intersection of Quantum Computing, Artificial Intelligence and Machine Learning

There are some obvious and not-so-obvious overlaps among various “advanced computing” concepts. Before I describe some of the inter-relationships among these concepts, it would be helpful to level-set the general definitions:

Classical Computing: is the form of data storage and analysis utilizing transistors in integrated circuits to turn switches on or off, hence storing a given computational state as a “bit”. These circuits are coordinated into logic gates to perform various instructions such as “AND”, “OR” and “NOT” and do so **in a sequential manner**. Today’s computers are increasingly fast and robust, having enjoyed Moore’s law for nearly 50 years. However, classical computers are beginning to hit an advancement ceiling and with the ever-increasing amount of data being collected and stored, the sequential nature of classical computing analysis is leading to longer and longer processing times for large data sets.

High-Performance Computing (HPC): is a technology that harnesses the power of supercomputers or computer clusters to solve complex problems requiring massive computation. While aggregating computing resources can improve overall power and speed, such increases in performance are linear (i.e., classical computing based), so an increasingly large set of resources is required as the data increases.

Quantum Computing: Quantum Computers (QCs) utilize evolving new technologies which take advantage of certain features of quantum mechanics. It uses “qubits” instead of classical computing bits and harnesses the properties of superposition, entanglement, and interference to perform calculations. Combining these quantum properties with a broader array of logic gates, QC’s can perform calculations simultaneously (instead of sequentially) and therefore much faster than classical computers. QCs are relatively new, and the existing devices are still not very powerful, but they are becoming more and more powerful all the time.

Artificial intelligence (AI): is intelligence demonstrated by machines, as opposed to natural intelligence displayed by animals including humans. In AI’s most basic form, computers are programmed to “mimic” human behavior using extensive data from past examples of similar behavior. AI applications include advanced web search engines (e.g., Google), recommendation systems (used by YouTube, Amazon, and Netflix), understanding human speech (e.g., Siri and Alexa), self-driving cars (e.g., Tesla), etc.

Machine Learning (ML): the study of computer algorithms that can improve automatically through experience and using data. It is seen as a part of artificial intelligence. Machine learning algorithms build a model based on sample data, known as training data, to make predictions or decisions without being explicitly programmed to do so. There are three types of machine learning: supervised learning (classification and regression), unsupervised learning (clustering and dimensionality reduction), and reinforcement (semi-supervised learning).

Big Data: refers to large, diverse sets of information that grow at ever-increasing rates. It encompasses the volume of information, the velocity or speed at which it is created and collected, and the variety or scope of the data points being covered (known as the "three v's" of big data). Data analysts look at the relationship between different types of data, such as demographic data and purchase history, to determine whether a correlation exists.

Quantum Machine Learning

Using these broad definitions, we can further refine this discussion to note that “Artificial Intelligence” today, is a general catch-all category for using classical computers to parse, analyze and draw conclusions. ML and Big Data are generally considered sub-sets of AI and HPC is a general catch-all for using mainframes, supercomputers and/or parallel processing to scale the power of classical computing. With the recent introduction of working QCs and given that QCs operate with different processes and logic, there is an evolving field known as “**Quantum Machine Learning**” (QML) at the intersection of these technologies.

Over the past few years, classical ML models have shown promise in tackling challenging scientific issues, leading to advancements in image processing for cancer detection, predicting extreme weather patterns, and detecting new exoplanets, among other achievements. With recent QC advances, the development of new *Quantum* ML models could have a profound impact on the world’s biggest problems, leading to breakthroughs in the areas of medicine, materials, sensing, and communications.

In a milestone discovery, IBM and MIT revealed the first experimental proof that the theory of combining quantum computing and machine learning could be achieved. They published their findings in Nature on March 13, 2019, using a two-qubit QC to demonstrate that QCs could bolster classification supervised learning.

TensorFlow, and PyTorch are leading platforms used for classical computing machine learning. TensorFLow is an end-to-end open-source platform with a comprehensive ecosystem of tools, libraries and resources that allow researchers and ML developers to easily build and deploy ML powered applications. PyTorch is also open source and has a machine learning library that specializes in tensor computations, automatic differentiation, and GPU acceleration.

TensorFlow Quantum

Reimagining these concepts for use on a QC, Google has released open sourced TensorFlow Quantum (TFQ) which provides quantum algorithm research and ML applications within the Python framework, designed to build QML models leveraging Google’s QC system. To build and train such models, users would do the following:

1. Prepare a quantum dataset
2. Evaluate a quantum neural network model
3. Sample or average measurements
4. Evaluate a classical neural networks model
5. Evaluate cost functions
6. Evaluate gradients and update parameters

Which is graphically depicted below:

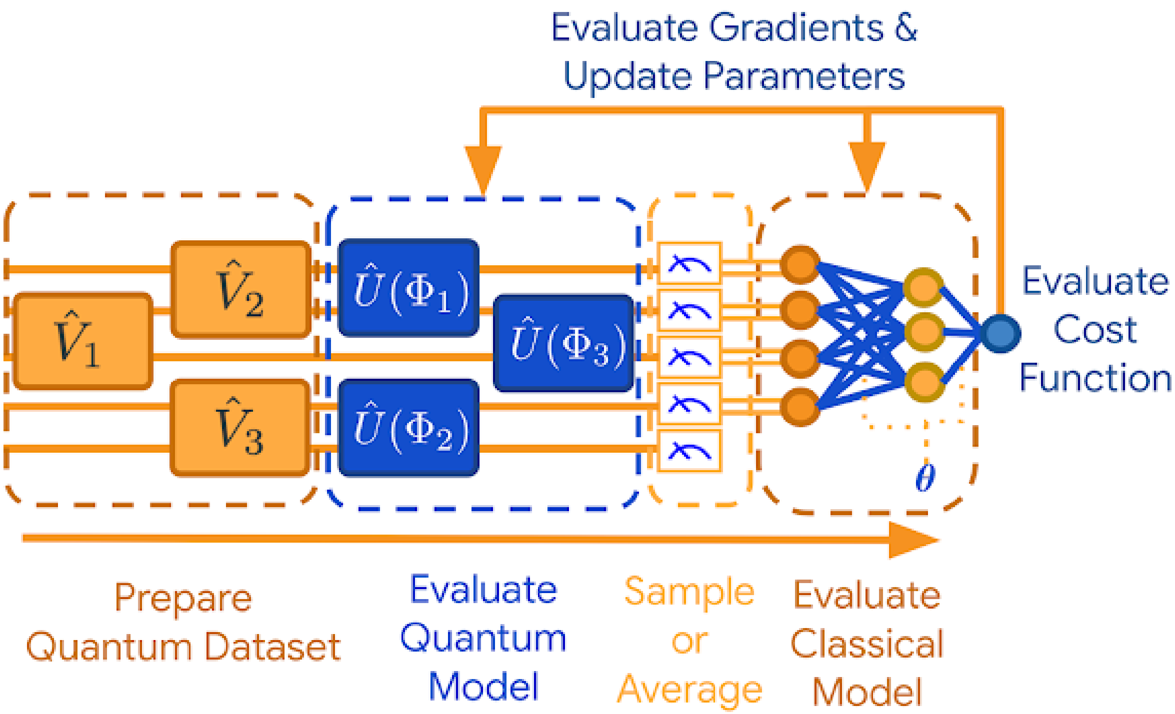

A key feature of TensorFlow Quantum is the ability to simultaneously train and execute many quantum circuits. This is achieved by TensorFlow's ability to parallelize computation across a cluster of computers, and the ability to simulate relatively large quantum circuits on multi-core computers.

PennyLane

Similarly, Xanadu's PennyLane is another open-source software framework for QML, built around the concept of quantum differentiable programming. It integrates classical ML libraries with quantum hardware and simulators, giving users the power to train quantum circuits. Companies such as Menten AI are using PennyLane to design novel drug molecules that can efficiently bind to a specific target of interest. Menten AI is seeking to develop new approaches that are beyond the reach of current classical computation by integrating QC and classical machine learning techniques.

PennyLane, is integrated with Amazon Braket, a fully managed quantum computing service from Amazon Web Services (AWS). Together with Amazon Braket, it seamlessly integrates classical machine learning (ML) libraries with quantum hardware and simulators, giving users the power to train quantum algorithms in the same way they train neural networks. Data scientists and machine learning researchers who work with TensorFlow or PyTorch on AWS will now have a way to experiment with quantum computing and see how easily it can fit into their workflows.

"Amazon Braket makes it easy for customers to experiment with quantum computing through secure, on-demand access to a variety of quantum hardware and fully managed simulators. We are delighted to be working with PennyLane to give our customers a powerful set of tools to apply proven and familiar machine learning concepts to quantum computing. Our goal is to accelerate innovation, and PennyLane on Amazon Braket makes it easy and intuitive to explore applications of hybrid quantum computing, an area of research that aims to maximize the potential of near-term quantum computing devices" said Eric Kessler, Sr. Product Manager for Amazon Braket.

Summary

While QC is still in its early stages, there are promising developments in applying QC to Artificial Intelligence/Machine Learning. Menten AI's use of this technology for drug discovery and Quantum Image Processing are but two examples of near-term applications. As the amount of stored data and images continues to explode, along with the increasing adoption of voice recognition tools (i.e., Alexa, Siri, etc.) utilization of QML will be vital to enabling efficient use of these evolving tools. I expect we'll see many more collaborations and tools in the QML space in the next few years.

Disclosure*: I have no beneficial positions in stocks discussed in this review, nor do I have any business relationship with any company mentioned in this post. I wrote this article myself and express it as my own opinion.*

References:

Uj, Anjaii, "Quantum Machine Learning: A Smart Convergence of Two Disruptive Technologies, " Analytics Insights, October 24, 2018

"What is Quantum Machine Learning," published by Discover Data Science, accessed February 20, 2022

Havlicek, Corcoles, Temme, Harrow, Kandela, Chow & Gambetta, "Supervised learning with quantum-enhanced feature spaces," *Nature*, March 13, 2019

Pennylane.ai, accessed February 20, 2022

TensorFlow.org, accessed February 21, 2022

Ho, Alan and Mohseni, Masoud, "Announcing TensorFlow Quantum: An Open Source Library for Quantum Machine Learning," Google AI Blog, March 9, 2020

"Menten AI Partners with Xanadu to Develop Quantum Machine Learning for Protein-Based Drug Discovery," PR Newswire, January 25, 2022.

If you enjoyed this post, please visit my website and enter your email to receive future posts and updates:
http://quantumleap.blog

Russ Fein is a venture investor with deep interests in Quantum Computing (QC). For more of his thoughts about QC please visit the link to the left. For more information about his firm, please visit Corporate Fuel. Russ can be reached at russ@quantumleap.blog.