Quantum Machine Learning: Recent Advances and Outlook

Wesley O'Quinn and Shiwen Mao

Abstract

Quantum computing is currently at the nexus of physics and engineering. Although current generation quantum processors are small and noisy, advancements are happening at an astounding rate. In addition, machine learning has played a crucial role in many recent advances. The combination of these two fields, Quantum Machine Learning, is a small but extremely promising new field with the possibility of unlimited abilities. This work seeks to provide an introduction to this emerging field, along with a discussion of recent advances as well as problems that are yet to be solved.

Introduction

Quantum computing is currently at the nexus of physics and engineering. This type of computing was primarily proposed by the physics community and, until recently, remained a vaguely theoretical concept. In spite of this, many well-known methods such as Shor's factoring, Grover's Search, and Linear Systems algorithms were formulated and promised paradigm shifting ability if practically realized. Although current generation quantum processors are small and noisy, advancements are happening at an astounding rate, due largely in part to government and private sector funding. Recently, the National Quantum Initiative Act was passed. This bill provided up to 1.2 billion dollars in research grant money to accelerate quantum related development. Private sector funding has also accelerated to provide startup capital and fund a variety of research.

One of the primary motivations for the development of quantum computers is the upcoming plateau of traditional computers. The exponential growth of transistors on a computer chip, as predicted by Moore's law, will soon come to a close. This is not due to any economic reason, but simply due to the laws of physics. Current generation transistors are roughly ten nanometers. It has been shown that transistors under seven nanometers begin to experience the effects of quantum tunnelling. This phenomenon is created when barriers in the transistor become arbitrarily small, that is, when the size of a gate reaches a certain thickness an electron can "jump" over the barrier, creating current where it should not be. This non-classical effect renders the transistors almost useless. Although chip manufacturers may be able to overcome this effect to some extent, the size of the transistor will basically reach its limit soon.

Note that a quantum computer is not made of smaller transistors, but rather quantum bits (or qubits) that harness the quantum effects that cause chaos in a classical system. Due to the effects of superposition, each qubit added is equivalent to doubling the power of the computer. This stands in stark contrast to needing double the number of transistors in order to double the processing power of a conventional computer.

On the other hand, machine learning has played a crucial role in many recent advances (e.g., wireless communication and networking systems). In particular, deep learning harnesses the power of extremely massive amounts of data. To properly utilize this information, more computing power is constantly necessary. Quantum computing represents the potential to properly utilize this data. The combination of these two fields, Quantum Machine Learning, is a promising new field with the possibility of extravagant abilities [1].

Before jumping into the topic at hand, it is helpful to provide a brief background of quantum computing. For the purposes of this article, we will define quantum computing as any process that utilizes the effects of quantum mechanics for improved computing capability. Most large computing corporations such as IBM, Google, and Microsoft, as well as smaller start-ups such as D Wave, Rigetti, and IonQ, have made great strides in developing the quantum hardware. To make development even easier, many of the companies mentioned provide access to their facility through the cloud free of charge. Although the type of hardware used varies significantly between different companies, the basic characteristics of the systems can be broken down into two separate types: Universal Quantum Computing and Quantum Annealing.

Universal Quantum Computing

A universal quantum computer uses three main properties to provide a fundamentally different type of computation. These three qualities are superposition, entanglement, and phase. Superposition basically means that a qubit can represent many more states than a traditional bit. For example, when a qubit is in superposition, it can be any number of states until made to collapse to either a one or a zero. This can be visualized in Fig. 1. Entanglement takes into account what Einstein called "spooky action at a distance." It basically means that gates can be connected via this property. Finally, phase cancellation or interference

Digital Object Identifier: 10.1109/MWC.001.1900341

The authors are with Auburn University.

is at the core of many quantum algorithms. This makes use of the fact that these qubits have not only magnitude, but also phase.

This type of computation has high potentials in machine learning. There is also reason to believe that with more research and development, other benefits may come to light. Currently, one of the key limitations is hardware development. Every processor in the field has major struggles with decoherence, which in essence destroys any information encoded in the qubits. This being the case, it is difficult to execute complex algorithms on such processors.

Quantum Annealing

Quantum annealing is a more limited type of computation mechanism, but has seen Moore's law type advancements over the past two years. This type of system finds a minimum of a certain type of function. The idea is to map problems onto this function and then use the quantum processor to solve the respective problem. To visualize this, the system can be thought of as a landscape of hills and valleys. The solution is tied to the minimum point. Just as water will end up in the lowest valleys, so too will the solution be marked by the annealer. It has been shown that the D Wave 2x system could outperform both simulated annealing and quantum monte carlo by up to a factor of 108 on certain optimization problems. More recent works have proved the practical applicability of the system [3].

Quantum Machine Learning

Once the basic hardware has been introduced, it is helpful to take a step back and analyze how these systems can be used from a machine learning perspective. As quantum computing is fundamentally different than traditional computing, so too can quantum machine learning be used in radically different fashions. For example, the fact that quantum annealers inherently find the lowest energy states means that optimization problems can be encoded into these systems. The quantum HHL algorithm — a method of solving linear systems of equations — which uses fundamental quantum effects, can be leveraged in many machine learning contexts. Grover's search algorithm takes advantage of the quantum property of amplitude amplification to mark solutions in an unsorted database. One of the implications of this is that clustering can potentially take place in a much faster manner. To summarize, quantum computing provides tools for machine learning developers that were not previously available. This allows for potentially faster processing of data as well as the ability to generate new forms of algorithms. With this in mind, it is important that one understands both current systems available as well as recent advances.

Motivations and Challenges

As mentioned in the introduction, a major motivation for quantum computers is the impending end of Moore's Law. But even more than that, quantum computing promises benefits that would not be realizable with classical hardware. For example, Shor's factoring algorithm is one of such first benefits. This particular algorithm will provide almost an exponential speedup with regard to factoring. This is proposed as a method for breaking RSA (Rivest-Shamir-Adleman) encryption. In addition, Grover's algorithm uses phase amplification to find items in an unsorted database. This could be used in a variety of applications from graph theory to machine learning. Finally, the quantum linear systems algorithm is particularly well suited for machine learning.

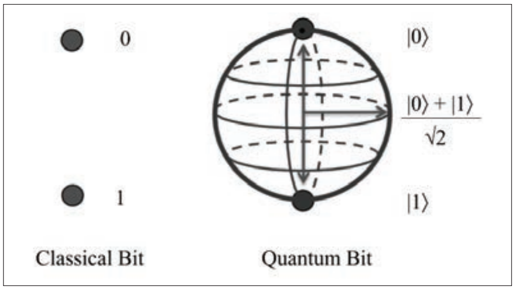

FIGURE 1. Superposition illustrated [2].

Even so, quantum computing faces massive engineering and programming issues. On the hardware side, noise plays a major role. Current generation quantum processors are limited almost solely to toy problems since very few operations can take place before noise degrades all the computations to a meaningless level. To limit this, processors must run at almost absolute zero temperature, greatly increasing the cost and creating many engineering problems. In addition, it is not clear how to fully harness the capabilities of the processor. Since it is so different from classical programming, algorithm development is slow and plagued with questions. These and other mainly engineering and programming issues must be addressed before the promises of quantum computing can be fully harvested.

Existing Platforms

Since there are a variety of methods to accomplish quantum computation, the hardware utilized varies greatly between companies. A qualitative comparison of the exiting platforms is given in Table 1. It is also helpful to note the information made available through [4], which is one of the few pieces of literature that provides extensive benchmarking data on a variety of platforms.

IBM Q

IBM Q, the quantum branch of IBM, is potentially the most well-known platform. Great care has been taken to tailor this system to the general public, which is available in an almost seamless manner through the cloud. It is not hard for a first time user to begin using the circuit composer, which allows one to design quantum circuits. Figure 2 provides a snapshot of Grover's algorithm implemented in the circuit composer.

Currently, a 14-qubit model is available to the general public, while their 20-qubit model can be used only by IBM Q clients. The hardware utilizes a transmon qubit. A microwave resonator is used to address and couple the qubits. The idea behind this system is to create electrically controlled solid-state quantum computers. Development can take place through Qiskit, which is an open-source quantum programming framework. The advantages of IBM Q is that a variety of materials are available intro-

Once the basic hardware has been introduced, it is helpful to take a step back and analyze how these systems can be used from a machine learning perspective. As quantum computing is fundamentally different than traditional computing, so too can quantum machine learning be used in radically different fashions.

System	Ease of Use	Available Online	Training Resources	Language	Pre-Built Functions	Hardware
IBM Q	✓	✓	✓	Python	×	32 qubit simulator + 14 qubit computer
IonQ	N/A	✓	×	N/A	N/A	Low error + 11 fully connected qubits
Regitti Computing	✓	✓	✓	Python	✓	16 qubits + integrated classic computation
D Wave	✓	✓	✓	Python	✓	Low Noise 2000 Qubit Quantum Annealer
Microsoft Quantum	✓	×	✓	Q#	✓	40+ qubit simulator
Google Quantum	N/A	×	✓	Python	✓	72 qubit processor

TABLE 1. Comparison of existing platforms.

ducing both the field of quantum computers and specifically IBM's version of quantum computers. The systems are available immediately for use anywhere, and the circuit composer is very intuitive. Furthermore, publications are available that utilized this system for research [5].

IonQ

IonQ has produced considerable results over the past few years, through the use of Ion trap technology. The qubits are housed in an ultra high vacuum chamber, and then precision lasers are used to connect the qubits and perform operations. The reason this company is unique is due to the low error scalability of their systems. See [4] for a hard data comparison of Trapped Ion technology with other platforms performed at the University of Maryland (UMD). Gaining access to the system entails an application process; only certain users can obtain access to the system. Moreover, literature on the software used is not readily accessible, especially to individuals new to the field, which could make any prospective project using this system more difficult.

Rigetti Computing

Rigetti's approach to the field is to develop a full stack solution. They plan to do this through a quantum-classical approach. Rather than viewing the quantum processor as a standalone unit, Rigetti seems to view it more as a piece of hardware such as a GPU. They have also recently released the Quantum Cloud Services (QCS). The core advantage of this system is the ability to run algorithms remotely through QCS. Both the classical and quantum hardware are available through servers. Any project done using this system will be easier on the software side, but may lack the quantum hardware that is available through other platforms.

D Wave

The D Wave system is highly different from the other systems. As mentioned, this particular computer is a type of quantum annealer. A helpful way to compare the two systems is to think of the annealer as an analog system, while the gate model is a digital system. This analogy will break down if pushed too far, but it is a helpful way to think of the process. This system basically finds the lowest energy state. If a problem can be mapped to the system, it will find a variety of possible optimal solutions. That said, the applications are still varied and useful. The interface and customer service of D Wave are second to none. The ability to run problems on the processor can be obtained quickly. A variety of functions have been made available, and the associated documentation is superb. In summary, D Wave is an excellent platform for more limited applications.

Other Platforms

Other major players in the field have not been mentioned, due to the fact that access to their hardware via the cloud is not currently available. Even so, some of these platforms provide meaningful resources. Microsoft has developed a language specifically for programming a quantum computer, referred to as Q#. They also provide a quantum development kit, which can integrate seamlessly with Azure to simulate 40+ qubits. Google's quantum branch operates under its artificial intelligence section. Currently, its quantum processor dubbed "Bristlecone" is extremely advanced, housing 72 qubits. One major contribution this team has made is the open source framework Cirq.

Recent Advances

Most advances in this field are very recent. This is true for two reasons: first, hardware and its availability through the cloud is an extremely recent event. Second, most of the research tends to build upon itself in a very foundational way. With the onset of quantum cloud computer systems, the potential for use is available in a way not seen in prior decades. Certain computational subroutines can be sent to cloud servers where they are processed. The information is then returned and can be utilized for a variety of applications. Frameworks should be developed that will be able to take full advantage of the system when it reaches maturity. If quantum computers hold to a Moore's law scheme, this situation will occur in the very near future. One of the ways this can be immediately utilized is in the field of machine learning.

Artificial Quantum Neural Networks

A recent publication in the field which drew significant attention [5] implemented an artificial neuron on an actual quantum processor. Tacchino *et al.* [5] demonstrated that a model based upon the classical Rosenblatt "perceptron" could be implemented on near-term hardware. The beauty of this model is that it can be trained by a hybrid quantum-classical scheme and shows exponential advantage in storage resources. This work puts forward that neural networks in particular meld almost seamlessly onto quantum hardware. This

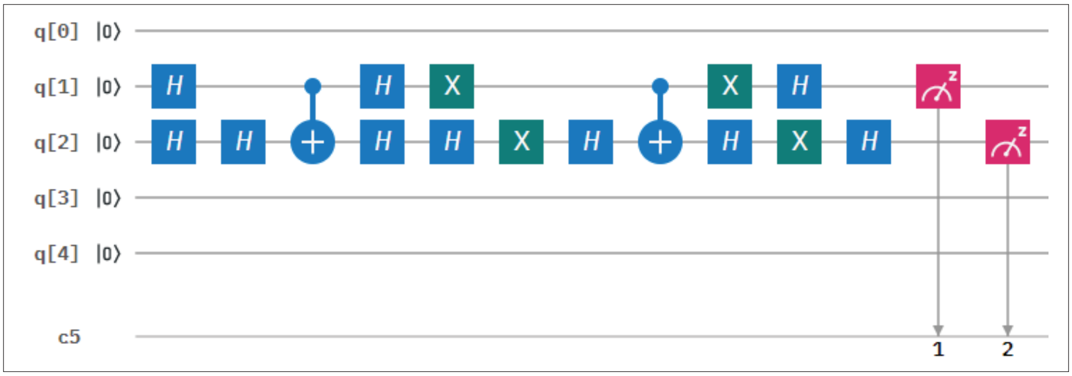

FIGURE 2. Grover's algorithm implementation using the circuit composer.

is true due to the fact that intrinsically, quantum mechanics has the property of representing and storing large complex valued vectors and matrices. Linear operations on such vectors can be performed as well. For simplistic representation, simple McCulloch-Pitts binary neurons were utilized. The model was then tested in a two-qubit model on an IBM Q quantum processor. A larger version was then implemented on the IBM quantum simulator. The results showed that the network could indeed be implemented, trained, and produce significant results. This work represents a compelling first step into the field of quantum neural networks.

Quantum Support Vector Machines

In a recent article produced by a collaboration of researchers at MIT and Google, a support vector machine was implemented with quantum hardware [6]. In this article, a method is given for implementing a quantum optimized binary classifier. It is interesting to note that for this system, where classical sampling algorithms would require polynomial time, exponential speedup can be achieved. The improvement was developed through a non-sparse matrix exponentiation technique, which efficiently achieved a matrix inversion of the training data inner-product matrix. This algorithm was developed specifically with Big Data classification in mind and proved that a quantum support vector machine (QSVM) could indeed be realized on hardware. In [7], a 4-qubit nuclear magnetic resonance test bench was used. In this study, the system was trained on standard character fonts and then used to classify handwritten characters. Although the test data was simplistic, the scheme performed admirably. The circuit utilized in this research can be seen in Fig. 3.

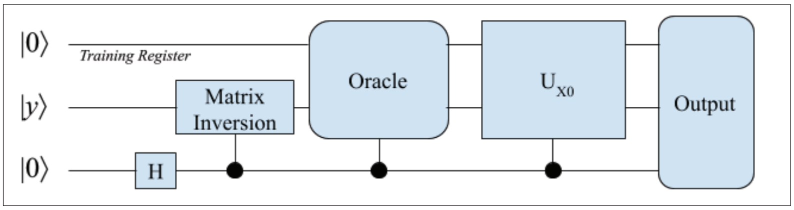

FIGURE 3. The QSVM circuit [7].

Learning Hidden Quantum Markov Models

Hidden Markov models (HMMs) are a special case of the Bayesian network family. This particular model is the premier algorithm used for speech recognition. It is also utilized in many other fields, including reinforced learning. Due to its widespread usage, it is important that the quantum applications of this algorithm be studied. This is especially relevant since the format of HMMs lends itself to smooth transition into the language of open quantum systems [8]. A variety of theoretical studies have been based on this concept. Monras *et al.* [9] introduced the idea of Hidden Quantum Markov Models (HQMM). This study was conducted by an in-depth look at the relationship between HMM and HQMM.

Recently, a collaborative group of researchers at the Georgia Institute of Technology and Carnegie Mellon University published a paper that accomplished three different tasks [10]. First, it was proven that HMM could be simulated on a quantum circuit. Next, the HQMM algorithm was reformulated to relax the constraints on quantum circuits. Finally, a learning algorithm is presented to estimate the parameters of an HQMM from data. This article provides many novel ideas that have yet to be implemented. Moreover, this work demonstrates that while the proposed HQMMs cannot model data any better than a sufficiently large HMM, it can better model the same data with fewer hidden states. An interesting characteristic of research related to quantum programming is that many times, advances can be made in classical computing through the work. The reason for this is the paradigm shift that must occur before progress can be made. The way quantum programming is done is fundamentally different from classical programming; therefore, ideas produced in the field tend to be novel in concept. For example, the algorithm produced in [10] can be utilized in both classical or quantum hardware.

Quantum Annealing

Along with the gate model type, annealing processors have grown extensively. For example, D Wave demonstrated the use of a 28-qubit system in 2007. In 2015, they announced that the 1000-qubit barrier had been broken. Finally, in 2017, a two thousand qubit model was released for commercial purchase. With promises for a lower noise and more connected 5000-qubit model to be released by mid 2020, the scalability is massive.

Traffic Optimization

Although not directly related to machine learning, Volkswagen's paper detailing the use of a quantum annealer provides meaningful insights into

As quantum machine learning is an emerging field, future research opportunities are highly evolving and almost endless. Interestingly, since theory advances faster than hardware, the most intriguing problems tend to be actually implementing the theoretical work on hardware.

high level computations with this system. In this paper [3], the team took into account 418 Beijing taxis with three possible routes. The target was to find the optimal route for each taxi. The problem was then mapped to the Quantum Unconstrained Binary Optimization problem. At its core, the system evolves probabilistically to the lowest energy level. This being the case, all problems must be formulated in such a way that this minimum energy state is the solution. Interesting results have been demonstrated in this project.

Restricted Boltzmann Machines

Quantum annealers lend themselves particularly well to probabilistic machine learning models. Inherently, quantum machines are probabilistic in nature. This is important since many problems are probabilistic. To use the forthright words of Richard Feynman, "Nature isn't classical ... If you want to make a simulation of nature, you'd better make it quantum mechanical." Restricted Boltzmann Machines seem to be a good match for these systems. Reference [11] proposes and implements such a scheme on actual hardware. In this paper, the results found through quantum annealing (QA) were compared with simulated annealing (SA). Under certain circumstances, the QA process returned the perspective ground state found by SA. Other times it returned one of the local minima. This seemed to arise from the lack of connectivity in the QA hardware. It will be interesting to see future results as new hardware becomes available. This seems to be one of the first works in machine learning based on QA.

Future Perspectives

As quantum machine learning is an emerging field, future research opportunities are highly evolving and almost endless. Interestingly, since theory advances faster than hardware, the most intriguing problems tend to be actually implementing the theoretical work on hardware. Before jumping into open problems, it would be helpful to mention a few papers that provide the background needed to approach these. Specifically, Nielsen and Chuang's textbook [12] is a very useful reference. In addition, Reference [13] furnishes an excellent introduction to the field of quantum machine learning. It discusses most of the work which had been done to date, as well as open problems and possible implementations.

Quantum Feedforward Deep Neural Networks

One interesting development moving forward will be the work built upon [5]. Very exciting paths of work are proposed in the discussion of this paper. The first is encoding continuously valued vectors rather than the binary model utilized. This would be equivalent to gray scale images rather than black and white. Even more interesting is the possibility of connecting multiple layers of quantum perceptrons together. The model could be fully implemented on quantum hardware and would effectively form a feedforward deep neural network.

Quantum Bayesian Networks

As mentioned in the review of [10], many novel theoretical ideas have been proposed that have yet to be implemented. Research in this area would be to provide experimental validation of the algorithms in question on a gate model system. One option could be based on the superconducting technology. Although noise is still a factor, a 16-qubit model is available through the cloud through IBM Q. This system would allow a more in-depth look into the problems in question. Also, it is interesting to note that all calculations utilized have been either synthetically generated or handwritten. If the system is implemented, the next step would be to utilize real world data. This novel research would represent the possibility for a wide range of wireless communications problems. Basically any process that utilizes HMM could potentially benefit from HQMM.

Graph Theory and Clustering

"Unsupervised learning (is) attractive in applications where data is cheap to obtain, but labels are either expensive or not available [14]." With the advent of Big Data, a great deal of unlabeled information is readily available for a variety of machine learning problems. But when the information is measured in the order of exabytes, training could be computationally expensive. Clustering is one of the most important tasks in unsupervised learning. To compound this, many clustering problems can be inherently represented as graph problems. In 2004, Durr *et al.* proposed a solution to the minimum spanning tree problem [15]. This was done through the utilization of a quantum enhanced version of Boruvka's algorithm. This specific algorithm was chosen over Kruskal's, because of its highly parallel nature. This work specifically shows that "if the connection matrix of the minimum spanning tree can be given by a quantum oracle, the computational time on a quantum computer can be reduced using Grover's algorithm to $O(N^{1.5})$." Since a minimum spanning tree can be easily turned into clusters by subtracting the k minus one connections, where k is the number of clusters, this algorithm lends itself naturally to clustering. The true impedance to this application (as well as most algorithms based on Grover's search) is the construction of the oracle. It is a problem that seems to defy logic, and is in itself an interesting problem to study.

Models Based on Quantum Annealing

Research in the area of quantum annealing seems to have been carried out almost solely by large corporations. This is probably due to the prohibitive cost of such machines. We predict this will change over the course of the next few years as cloud models are introduced. Quantum annealers are especially enticing due to the demonstrated scalability. Also due to the software utilized, they are more simplistic to develop algorithms for use. Finally, real world benefits have already been proven through traffic routing problems as well as satellite routing [15].

Conclusion

To use the words of Chad Rigetti, "Quantum computing is arguably the most sophisticated technology that humans have ever developed." Although many would agree with this statement, the field is still in its infancy, while the promises of quantum computing are almost endless. Therefore, research in this field is of the utmost value. This article constitutes a literature review of existing platforms and important issues and hopefully will provide resources for those interested in pursuing a deeper understanding and a desire to apply quantum machine learning to their respective fields.

Acknowledgments

This work is supported in part by the NSF under Grant ECCS-1923717 and ACI-1659845, and by the Wireless Engineering Research and Education Center (WEREC) at Auburn University. We would like to thank Dr. Peter Mueller with the IBM Zurich Research Laboratory for his suggestions with regard to the IBM Q systems, and the authors of [4] for constructive feedback. We would also like to thank Dr. Mark Novotny with Mississippi State University for his explanation of his recent work with quantum annealers. We are indebted to Mr. Michal Stechly with Zapata Computing for his exceptional work in explaining concepts related to NISQ processors and related suggestions. Finally, Dr. Joel Gottlieb at D-Wave Systems has been an invaluable resource in both this work and future works.

References

[1] V. Dunjko, J. M. Taylor, and H. J. Briegel, "Quantum-Enhanced Machine Learning," *Phys. Rev. Lett*., vol. 117, p. 130501, Sept. 2016; available: https://link.aps.org/doi/10.1103/PhysRevLett.117.130501.

[2] A. Daley, "Quantum Optics and Quantum Many-Body Systems: Quantum Computing," Jan. 2012; available: http://qoqms.phys.strath.ac.uk/research qc.html.

[3] F. Neukart *et al*., "Traffic Flow Optimization Using a Quantum Annealer," Aug. 2017; available: https://arxiv.org/abs/1708.01625.

[4] P. Murali *et al*., "Full-stack, Real-system Quantum Computer Studies: Architectural Comparisons and Design Insights," *Proc. ISCA'19*, Phoenix, AZ, June 2019, pp. 1–14.

[5] F. Tacchino *et al*., "An Artificial Neuron Implemented on an Actual Quantum Processor," *Nature npj Quantum Information*, vol. 5, no. 26, Mar. 2019.

[6] P. Rebentrost, M. Mohseni, and S. Lloyd, "Quantum Support Vector Machine for Big Data Classification," *Phys. Rev. Lett*., vol. 113, no. 13, Sept. 2014, p. 130503.

[7] Z. Li *et al*., "Experimental Realization of a Quantum Support Vector Machine," *Phys. Rev. Lett*., vol. 114, no. 14, Apr. 2015, p. 140504.

[8] L. A. Clark *et al*., "Hidden Quantum Markov Models and Open Quantum Systems with Instantaneous Feedback," arXiv e-prints, June 2014, p. arXiv:1406.5847; available: https://arxiv.org/abs/1406.5847.

[9] A. Monras, A. Beige, and K. Wiesner, "Hidden Quantum Markov Models and Non-Adaptive Read-Out of Many-Body States," arXiv e-prints, Feb. 2010, p. arXiv:1002.2337; available: https://arxiv.org/abs/1002.2337.

[10] S. Srinivasan, G. Gordon, and B. Boots, "Learning Hidden Quantum Markov Models," arXiv e-prints, Oct. 2017, p. arXiv:1710.09016; available: https://arxiv.org/abs/1710.09016.

[11] Y. Koshka *et al*., "Determination of the Lowest-Energy States for the Model Distribution of Trained Restricted Boltzmann Machines Using a 1000 Qubit D-Wave 2x Quantum Computer," *Neural Computation*, vol. 29, no. 7, July 2017, pp. 1815–37.

[12] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, New York, NY: Cambridge University Press, 2000.

[13] J. Adcock *et al*., "Advances in Quantum Machine Learning," arXiv e-prints, Dec. 2015, p. arXiv:1512.02900; available: https://arxiv.org/abs/1512.02900.

[14] J. Biamonte *et al*., "Quantum Machine Learning," *Nature*, vol. 549, no. 7671, Sept. 2017, pp. 195–202.

[15] C. Durr *et al*., "Quantum Query Complexity of Some Graph Problems," arXiv e-prints, Jan. 2004, pp. quant–ph/0 401 091; available: https://arxiv.org/abs/quant-ph/0401091.

Biographies

Wesley O'Quinn received the B.S. degree in electrical engineering and B.A. degree in Chinese language from Georgia Southern University, Statesboro, GA as an honors student, both in 2019. He is with the Southern Nuclear Operating Company as an engineer and is pursuing a Ph.D. in electrical and computer engineering at Auburn University, Auburn, AL. His research interests include machine learning/deep learning, quantum computing, signal processing, wireless system design, and energy harvesting.

Shiwen Mao [S'99, M'04, SM'09, F'19] received a Ph.D. in electrical and computer engineering from Polytechnic University, Brooklyn, N.Y. in 2004. He is the Samuel Ginn Professor and Director of the Wireless Engineering Research and Education Center at Auburn University, Auburn, AL. His research interests include wireless networks, multimedia communications, and smart grid. He is a co-recipient of several best paper awards, several service awards, and a best demo award from the IEEE Communications Society (ComSoc). He received the IEEE ComSoc TC-CSR Distinguished Technical Achievement Award in 2019, the NSF CAREER Award in 2010, and The 2004 IEEE Communications Society Leonard G. Abraham Prize in the Field of Communications Systems. He is a Fellow of the IEEE.

To use the words of Chad Rigetti, "Quantum computing is arguably the most sophisticated technology that humans have ever developed." Although many would agree with this statement, the field is still in its infancy, while the promises of quantum computing are almost endless. Therefore, research in this field is of the utmost value.