Quantum: A New Kind of Knowledge Discovery

Ramin Ayanzadeh[1,2] and Milton Halem[2]
[1] Georgia Institute of Technology
[2] University of Maryland, Baltimore County

The New Game in Town

While the first solid-state device (known as transistor) was being developed at Bell lab in the mid-twentieth to replace vacuum-tubes [30], artificial intelligence (AI) was being conceptualized by a generation of scientists, mathematicians, and philosophers. In 1950, Alan Turing suggested two criteria for machine intelligence: memory for enabling machines to store and retrieve data, and reasoning (i.e., having the capacity to process data) [31]. Since then, trends in doubling the transistor count, characterized by Moore's law, have catalyzed AI advancements. Nowadays, AI applications have access to not only large-scale memories but also high-performance computing (HPC) resources.

After decades of predomination, the era of Moore's law is drawing to a close. Are we prepared for the end of Moore's law? Can digital systems keep pace with ever-increasing demand for data storage and information processing capacity? The microelectronics industry (will be known as the nanoelectronics industry in the near-term) trying to identify new materials and devices to replace the 50 years old transistor technology—including, but not limited to, non-classical CMOS (such as new channel materials), and alternatives to CMOS (e.g., spintronics, single-electron devices, and molecular computing) [30]. Although the microelectronics industry will continue to reduce electronic devices' costs, there are theoretical boundaries that can limit future micro/nano-electronic processing devices' computing power. However, it is worth noting that AI applications can still benefit from the dramatic increase in memory systems' speed and storage capacity.

Transitioning from vector-based computation, in central processing units (CPUs), to matrix-based computation resulted in the emergence of graphical processing units (GPUs) that have reshaped the landscape for accelerated computing, namely in realms of scientific computing, high-performance computing, machine learning, and big data analytics. In the same manner, custom design application-specific integrated circuits (ASICs) for extending computations (from vectors and matrices) to tensors (i.e., complex and higher-order objects), can repeat providing a disruptive capability. As an example, the tensor processing unit (TPU), by Google, is an accelerator for near-real-time deep learning applications with low latency that has demonstrated throughput improvements of over 15–30 and power efficiency improvement of 30–70 over current CPUs and Kepler generation of GPUs, albeit lower precision computations [17]. The supercomputing community tries to address the limits in shrinking transistors' size through parallelism (i.e., distributing data and processes), which has emerged a new kind of race (or war) between the United States, China, and perhaps other players like Russia. Besides national security and economic concerns in developing the next generation of supercomputers, the near-term future of AI can mainly depend on the result of the supercomputing race. Hence, we can expect that the winner of the supercomputing game will achieve AI supremacy.

Post-Moore Era: Emerging New Technologies

Besides all aforementioned trends in advancing the computing power of the next generation of accelerators, there is a growing consensus that we will eventually need different type(s) of computing machinery. In post-Moore era, therefore, we explore non-Von-Neumann architectures (such as zero-instruction set and single-instruction set computers) and non-digital systems (namely quantum computers) for emerging next-generation of accelerators. As an example, neural processing units (NPUs) or neuromorphic chips—namely the Neurosynaptic System by IBM, the SpiNNaker System by the University of Manchester, Intel's Loihi chip, and memristors based systems—are zero-instruction set computers inspired by the human brain that have demonstrated a dramatic speedup in implementing AI models, more specifically deep neural networks.

In 1982, Richard Feynman proposed the use of quantum systems for simulating quantum processes [16]. But why does one need a quantum computer to simulate quantum processes when our classical silicon-based computers can solve all types of physics problems from relativistic to Newtonian models. The answer lies in the following three seemingly unreal properties that atoms exhibit when cooled down to near absolute zero, i.e., below 20 milli-Kelvin that have come to be accepted today. First, the property of superposition. While a classical bit can only be either 0 or 1, a quantum bit (qubit) is a two-level quantum system that can be 0 and 1 simultaneously (superposition of 0 and 1) with corresponding probabilities. Let's say one were to encode information into a quantum state or analogously into a quantum bit, whether the state is created by magnetic fields about a gallium arsenide chip or by trapping ions with lasers or forming arrays of cesium atoms at one-millionth of a degree, these states represented by a 1 or 0 would form the bit structure of a quantum computer. A quantum register with n qubits can simultaneously be in 2^n arrangements. This means a quantum system with 100 qubits can have 2^{100} arrangements, more arrangements than atoms in the universe. This would enable unbelievable searches. The second property known as entanglement is even stranger. Assume two such sets of quantum bits were correlated here in Washington DC, and then one set of quantum bits is separated by the Pacific ocean in a site in China. Then any operation on the quantum bits here in Washington DC would simultaneously enact at the entangled quantum bits in China. Finally, the third property is known as interference. Unlike classical wave functions that can only interfere with each other, in quantum mechanics, an individual particle can cross its own trajectory and interfere with itself. Quantum interference enables us to bias the measurement of a quantum register toward the desired outcome.

These three steps enable quantum operations to perform functions that would be near impossible with classical computers. The challenge to effect such a quantum computer is that these states are so unstable to minor temperature, radiation, and vibrational effects such that noise limits the time to conduct operations on the quantum bits. The interest in exploring the potentials of this quantum technology was ignited when Peter Shor demonstrated that the factoring problem is tractable or amenable in the realm of quantum computing [29]. Introducing the first programmable quantum processing unit (QPU) by D-Wave systems in 2011 [20], although it was not a universal quantum computer, revealed that quantum technology is still feasible in the near-term. We are confident that quantum science has the potential to emerge as a transforming technology, and quantum technology promises new capabilities in demonstrating quantum advantage in several domains, ranging from problem-solving (known as quantum computing) to sensing, communication, and simulation of quantum systems.

Quantum computing is a different information processing paradigm that takes advantage of quantum mechanics for addressing computationally intensive problems that are intractable in the realm of classical (or digital) computing. There are several models for the realization of quantum information processing. Although these models are theoretically equivalent, their underlying concepts, as well as realization requirements, are significantly different from each other. Circuit models in quantum computing are closely analogous to classical computers in which a sequence of quantum operations are applied to quantum registers that follows the control flow of a quantum algorithm to process the quantum information [24]. Considering the fact that quantum computers are fundamentally different machinery, running classical algorithms on a quantum processor are not plausible, and (for every problem) we need to design a new quantum algorithm [3, 7].

Adiabatic quantum computers rely on *Quantum Adiabatic Evolution* for quantum information processing. According to the *Adiabatic Theorem*, adiabatic quantum computers are polynomially equivalent to the circuit models. Unlike previous models in quantum information processing, qubits of adiabatic quantum computers do not perform discrete operations. Indeed, adiabatic quantum computers receive a Hamiltonian (also called energy function) as input whose ground state represents the solution for the problem that we are trying to solve. Afterward, qubits are adiabatically evolved from some initial unknown state to a final state that minimizes the energy function [15].

While there are several competing approaches for the physical realization of circuit model quantum computers (namely superconductors, trapped ions and cold atoms), adiabatic quantum computers have not been physically implemented yet. The perspective, however, has resulted in introducing a novel optimization method called quantum annealing [21], and consequently emerging the quantum annealers as a single-instruction set computer, a.k.a. Ising processing units (IPUs) [20]. From a computing viewpoint, IPUs—namely digital annealers by Fujitsu (based on FPGA and recently CMOS technologies), optical Ising machines (OIMs) at the University of Rome La Sapienza (based on photons), and D-Wave quantum annealers (based on superconducting qubits)—can only sample from the ground state (i.e., a configuration with the lowest energy value) of a given (quadratic) objective function, known as Ising model. In the Ising model (a.k.a. spin-glass model), pairwise correlations between interacting variables (or spins) can model complex probability distributions or represent computationally hard problems [13]. While quantum annealers are not fully resemblance of adiabatic quantum computers [2, 4], recent studies have revealed the potential of exploring quantum advantage in the realm of quantum annealers [27, 11, 25, 8, 10, 6, 5, 23, 22]. Moreover, by sampling from high-dimensional probability distributions, one can leverage quantum annealers for many applications in artificial intelligence, machine learning and signal processing [11, 1]. It is worth highlighting that there are other proposals for quantum computing, such as measurement-based or one-way quantum computing, and we are in the era of exploring new emerging computing technologies.

NISQ Era: New Discoveries

Fault-tolerance quantum computing relies on continuous error correction; nevertheless, existing and near-term quantum computers cannot fully accommodate quantum error-correction techniques. Hence, until we can bypass several technological barriers, we are limited to explore noisy intermediate-scale quantum (NISQ) computers for exploring the quantum advantage [28]. One may say that we need thousands (and even millions) of physical qubits to achieve the quantum advantage, and we should not expect to be able to run any quantum algorithm on a near-term quantum computer. While this attitude sounds limiting, we should highlight quantum computers are not exclusive for only running quantum algorithms for addressing certain types of (classically intractable) problems. Indeed, quantum computing provides a different platform (and viewpoint) for addressing challenging applications. For example, calculating the greatest common divisor (GCD) of two numbers can require us to find all factors of two given numbers, and we know that factoring is a non-trivial task. However, in the realm of modular arithmetic, using the Euclidean algorithm, finding GCD is a trivial problem. In the same manner, when we put on our quantum computing lenses, we can see that finding the period of a given function is tractable; therefore, we can address the problem of prime factoring by applying Shor's algorithm [29].

From an application point of view, we can be optimistic about exploring quantum advantage through applying NISQ computers on (at least) two broad classes of problems: optimization and simulation. In this sense, optimization can be the first candidate to demonstrate the quantum advantage in the NISQ era. In the circuit model of quantum computing, variational quantum algorithms (a.k.a. classical-quantum hybrid schemes) such as quantum approximate optimization algorithm (QAOA) and variational quantum eigensolver (VQE) can address optimization applications that are intractable in the realm of classical computing [14]. In fact, we can expect the next generation of NISQ processors (namely cold atoms) to be capable of executing large and deep enough quantum circuits, which is the bottleneck for demonstrating the quantum advantage using QAOA and VQE [18]. Similarly, we can look at quantum annealers as a NISQ device, but for the adiabatic model of quantum computing. Considering the fact that quantum annealers are easier to scale, although their capacity is limited, we can expect to see the quantum advantage in the realm of quantum annealers [22]. Physicists are also excited about using NISQ devices for exploring the physics of many entangled particles. We know that quantum computers can simulate any natural process, and NISQ devices provide a valuable platform for discovering new aspects of physical processes [28]. To this end, we can expect new discoveries when we had access to NISQ computers with a few hundred qubits.

While most current quantum artificial intelligence studies propose applying quantum accelerators to hard AI problems, NISQ devices are less likely to provide a disruptive capability. Indeed, in NISQ era, we can expect AI to improve the fidelity of near-term quantum computers. As an example, several studies have reported notable improvements in the use of machine learning techniques (e.g., reinforcement learning) in not only finding better solutions but also in making the QAOA results reproducible. Recent studies have suggested that machine learning models can mitigate the measurement error, which is the most error-prone operation in most current quantum circuits. Since NISQ devices are susceptible to various error sources, a large number of trials are needed, and the correct answer is inferred based on the distribution of outcomes. In this context, deep networks can be trained to learn and mitigate the system noise of NISQ devices.

In the realm of adiabatic quantum computing (and quantum annealing), Boolean satisfiability (SAT) has been suggested as an intermediate problem representation for casting classical algorithms (implemented in classical programming languages) to Ising models, executable by quantum annealers and IPUs [3, 7, 6, 7]. Machine learning techniques, namely reinforcement learning, have the potential to notably improve the performance and robustness of the quantum annealers [9].

We may envision that quantum computers in the future will be flown in space where the gravity-free environment will help arrays of Bose-Einstein condensates of atoms (i.e., a state of matter where sub-atomic particles, cooled to very near absolute zero, coalesce into a single entity on a macroscopic scale), the qubits of the future quantum computer that can maintain longer coherent times and enable easier transitions to B-E states when increased to tens of thousands or even millions of qubits. This conjecture was successfully established by the flight of the ColdQuanta Lab to the International Space Station in 2019 by JPL to demonstrate the B-E state. As a by-product, it was found that the gravity free environment also increased coherent times in space. Thus the most beneficial quantum computing applications will occur from space where not only will quantum cyber information deal with cybersecurity, but these superquantum computers will be used to maintain the scheduling of the global internet, the national power grids, 100G communication in the new era, global tracking of all shipped products both physical and digital, and the management of global climate warming mitigation.

We are in an opportune time (the era of new discoveries) to employ the quantum computing lenses for not only finding new problems but also trying to solve classically intractable problems.

Post-NISQ Era: Quantum Advantage

We are in the NISQ era where we are limited to small and noisy quantum processors, but we are confident that fault-tolerance quantum computing is on the horizon. Will we be able to immediately demonstrate quantum advantage after having a reliable/useful quantum computer? Perhaps we will be able to start running quantum algorithms and achieve the promised speedup. Shor's algorithm is the first example that comes up; however, the potential of quantum science and technology is beyond the capability of currently known famous quantum algorithms. Recent studies suggest that quantum computing will be a disruptive technology for all scientific disciplines, and it will revolutionize biotechnology, quantum chemistry and material science.

Quantum artificial intelligence (QAI) and quantum machine learning (QML) are emerging fields that aim to leverage quantum computing for addressing certain types of problems that are intractable in the realm of classical computing. Although we do not expect NISQ accelerators to be the game-changer in the realm of artificial intelligence, we can expect fault-tolerance quantum computers to be a turning point. Can quantum computing revolutionize AI and machine learning? While most studies are optimistic about leveraging quantum computers to accelerate AI and machine learning applications, we should not expect quantum machine learning to fully outperform classical machine learning [19]. There are theoretical boundaries that limit the potential advantage of quantum machine learning over classical machine learning when we are interested in improving the average-case error. On the other hand, quantum machine learning can outperform classical models of learning when we are concerned about the worst-case scenarios [19]. Thus, the performance of quantum AI and quantum machine learning will mostly depend on applications rather than the model, so we need to find problems where worst-case analysis matters. Weather prediction can be an example where we can expect quantum machine learning to be a transforming technology. Similarly, we may expect quantum-assisted predictive models to play a crucial role in finance applications.

In the biopharma industry, developing a drug product from the initial scientific hypothesis can take more than ten years and billions of dollars before being commercialized. The vast majority of the time, effort, and cost are spent on experimental design and characterizations. Despite considerable progress in classical computers for modeling macromolecules such as proteins, ligands, and peptides, many molecular biology and biophysics challenges remain computationally infeasible. Numerically calculating the full electronic wavelength of a drug product is expected to take longer than universe age, even on the current most powerful supercomputers. Current well-designed algorithms for predicting protein structure and folded state of a native protein and calculating the binding affinity of a ligand for a macromolecule require computational resources beyond the capacity of the most available supercomputers. For example, for drug discovery, predicting the structure of a protein is essential. A protein structure can experimentally be evaluated by X-ray, NMR, and a more recent powerful Cryo-TEM technique. However, all these techniques have a lot of noise and do not provide an exact image of the protein. Powerful modeling can be utilized to prevent time-consuming efforts of experiments. It needs precise modeling to include different parameters (presence of water molecules, solvent, electrostatic interactions, hydrophobicity, etc.) as protein is a flexible macromolecule, making this process highly challenging. Finding the global minimum energy of a protein, based on its energy landscape, is another interesting topic in the field as characteristic, effectiveness, toxicity, pharmacokinetics, etc., of the protein, depends on its free energy. The structure which is correspondent to the minimum energy is the most effective and stable structure. Another approach in the biopharma industry is the protein-ligand docking technique utilized to predict whether a ligand binds to a protein. This technique is based on scoring the predicted model protein based on previously available data for proteins that have similarities to the protein of interest. Docking algorithm is also computationally challenging as it requires incorporating macromolecule flexibility and the effect of water molecules and solvents around it [26, 12]. Quantum computing promises to speed up drug discoveries exponentially, and we are optimistic that machine learning can expedite the achievement of quantum advantage. It is worth highlighting that similar concepts can be applied in material science to predict characteristics of new unknown materials.

We can (or should) imagine a different world in the era of quantum knowledge discovery. While quantum computing promises to jeopardize the security of current cryptographic systems, we should have enough time to come up with novel and quantum-resistant systems (i.e., post-quantum cryptographic systems) prior to the emergence of fault-tolerance quantum computers. We should expect quantum technology (which is not limited to only quantum computing) to significantly impact the vast range of industries, specifically the healthcare sector. In the era of quantum knowledge discovery, the biopharmaceutical industry (as an example) will not be considered as a (very) high-risk sector, and we can expect novel discoveries to flourish. Besides its economic impacts, we can expect a notable improvement in life expectancy through new drug discoveries and novel (early) diagnosis systems. Quantum science and technology can also enable us to predict (and prepare for) the upcoming climate changes that can be worse than the current pandemic. And finally, unfortunately, quantum technology can increase the gap between developed and underdeveloped countries from the knowledge discovery perspective. In the era of quantum knowledge discovery, access to premium quantum resources will play a crucial role in scientific research studies. Hence, winner(s) of the quantum race will lead the world in producing both science and scientists.

References

[1] Steven H Adachi and Maxwell P Henderson. Application of quantum annealing to training of deep neural networks. *arXiv preprint arXiv:1510.06356*, 2015.

[2] Tameem Albash and Daniel A Lidar. Adiabatic quantum computation. *Reviews of Modern Physics*, 90(1):015002, 2018.

[3] Ramin Ayanzadeh. *Leveraging Artificial Intelligence to Advance Problem-Solving with Quantum Annealers*. PhD thesis, University of Maryland, Baltimore County, 2020.

[4] Ramin Ayanzadeh, John Dorband, Milton Halem, and Tim Finin. Multi-qubit correction for quantum annealers. *Scientific Reports*, 11(1):1–12, 2021.

[5] Ramin Ayanzadeh, John Dorband, Milton Halem, and Tim Finin. Quantum-assisted greedy algorithms. In *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*, pages 4911–4914. IEEE, 2022.

[6] Ramin Ayanzadeh, Milton Halem, and Tim Finin. Solving hard sat instances with adiabatic quantum computers. In *AGU Fall Meeting Abstracts*, volume 2018, pages IN41B–27, 2018.

[7] Ramin Ayanzadeh, Milton Halem, and Tim Finin. Sat-based compressive sensing. *arXiv preprint arXiv:1903.03650*, 2019.

[8] Ramin Ayanzadeh, Milton Halem, and Tim Finin. An ensemble approach for compressive sensing with quantum annealers. In *IGARSS 2020-2020 IEEE International Geoscience and Remote Sensing Symposium*, pages 3517–3520. IEEE, 2020.

[9] Ramin Ayanzadeh, Milton Halem, and Tim Finin. Reinforcement quantum annealing: A hybrid quantum learning automata. *Scientific reports*, 10(1):1–11, 2020.

[10] Ramin Ayanzadeh, Seyedahmad Mousavi, Milton Halem, and Tim Finin. Quantum annealing based binary compressive sensing with matrix uncertainty. *arXiv preprint arXiv:1901.00088*, 2019.

[11] Rupak Biswas, Zhang Jiang, Kostya Kechezhi, Sergey Knysh, Salvatore Mandra, Bryan O'Gorman, Alejandro Perdomo-Ortiz, Andre Petukhov, John Realpe-Gómez, Eleanor Rieffel, et al. A nasa perspective on quantum computing: Opportunities and challenges. *Parallel Computing*, 64:81–98, 2017.

[12] Yudong Cao, Jhonathan Romero, and Alán Aspuru-Guzik. Potential of quantum computing for drug discovery. *IBM Journal of Research and Development*, 62(6):6–1, 2018.

[13] Arnab Das and Bikas K Chakrabarti. Colloquium: Quantum annealing and analog quantum computation. *Reviews of Modern Physics*, 80(3):1061, 2008.

[14] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.

[15] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *arXiv preprint quant-ph/0001106*, 2000.

[16] Richard P Feynman. Simulating physics with computers. *Int. J. Theor. Phys*, 21(6/7), 1982.

[17] Kathleen E Hamilton, Catherine D Schuman, Steven R Young, Ryan S Bennink, Neena Imam, and Travis S Humble. Accelerating scientific computing in the post-moore's era. *ACM Transactions on Parallel Computing (TOPC)*, 7(1):1–31, 2020.

[18] Matthew P Harrigan, Kevin J Sung, Matthew Neeley, Kevin J Satzinger, Frank Arute, Kunal Arya, Juan Atalaya, Joseph C Bardin, Rami Barends, Sergio Boixo, et al. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics*, pages 1–5, 2021.

[19] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Information-theoretic bounds on quantum advantage in machine learning. *arXiv preprint arXiv:2101.02464*, 2021.

[20] Mark W Johnson, Mohammad HS Amin, Suzanne Gildert, Trevor Lanting, Firas Hamze, Neil Dickson, R Harris, Andrew J Berkley, Jan Johansson, Paul Bunyk, et al. Quantum annealing with manufactured spins. *Nature*, 473(7346):194, 2011.

[21] Tadashi Kadowaki and Hidetoshi Nishimori. Quantum annealing in the transverse ising model. *Physical Review E*, 58(5):5355, 1998.

[22] A.D. King, J. Raymond, T. Lanting, et al. Scaling advantage over path-integral monte carlo in quantum simulation of geometrically frustrated magnets. *Nature Communications*, 12(1113), 2021.

[23] Ohzeki Masayuki. Breaking limitation of quantum annealer in solving optimization problems under constraints. *Scientific Reports (Nature Publisher Group)*, 10(1), 2020.

[24] Michael A Nielsen and Isaac L Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2010.

[25] Daniel O'Malley, Velimir V Vesselinov, Boian S Alexandrov, and Ludmil B Alexandrov. Nonnegative/binary matrix factorization with a d-wave quantum annealer. *PloS one*, 13(12):e0206653, 2018.

[26] Carlos Outeiral, Martin Strahm, Jiye Shi, Garrett M Morris, Simon C Benjamin, and Charlotte M Deane. The prospects of quantum computing in computational molecular biology. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 11(1):e1481, 2021.

[27] Alejandro Perdomo-Ortiz, Neil Dickson, Marshall Drew-Brook, Geordie Rose, and Alán Aspuru-Guzik. Finding low-energy conformations of lattice protein models by quantum annealing. *Scientific reports*, 2:571, 2012.

[28] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

[29] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

[30] Scott E Thompson and Srivatsan Parthasarathy. Moore's law: the future of si microelectronics. *Materials today*, 9(6):20–25, 2006.

[31] Alan M Turing. Computing machinery and intelligence. In *Parsing the turing test*, pages 23–65. Springer, 2009.