Horizon Scanning Series

The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing

Quantum Machine Learning

This input paper was prepared by Lloyd Hollenberg

Suggested Citation

Hollenberg, L (2018). Quantum Machine Learning. Input paper for the Horizon Scanning Project “The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing” on behalf of the Australian Council of Learned Academies, www.acola.org.

Quantum Machine Learning

L. Hollenberg, School of Physics, University of Melbourne

Intro:
Quantum computer technology and software are in a phase of rapid development. A key area of interest is the confluence of quantum computing (or quantum technology more broadly) and machine learning (ML) – i.e. quantum machine learning (QML) [1,2]. There are several themes emerging in QML: quantum algorithmic versions of linear algebra tasks required or associated with ML; using specialised quantum hardware to assist ML; development of fully quantum neural networks; the use of classical ML for the development of quantum technology itself. At present, given that these are all relatively nascent areas, the nature of the impact of quantum technology on ML is largely an open question. However, it is clear that the confluence of quantum information and machine learning is an exciting area and will continue to attract a lot of interest – increasing the level of interaction between the communities is key to shaping and capturing future applications.

Quantum computers:
The development of quantum computers, on both experimental and theoretical fronts, has accelerated in the past few years as industry and governments have increased their investment into quantum technology [3]. The fundamental component in a quantum computer is the quantum bit, or qubit. Qubits can be formed from a range of physical systems which have distinct controllable states, exhibit quantum behaviour, and can be read-out: e.g. superconducting circuits, photonic systems, trapped atoms, and spins in semiconductors, to name a few [4]. In contrast to classical bits, a qubit can be a quantum superposition of 0 and 1 states prior to measurement. In a quantum computer binary strings can be encoded over multiple qubits, and the subsequent quantum register put into a superposition of states. Contrary to popular belief, the power of a quantum computer is not solely derived from the superposition over binary numbers, or inherently parallel evaluation of functions, rather it is the ability to interact the qubits and perform quantum logic generating entanglement and non-classical interference effects that effectively modify the probability of measuring certain binary outcomes (i.e. the “answer”). Put another way – a classical machine will search for the needle in the haystack by sifting through all hay stalks, but a quantum computer does it very differently: not by searching all hay stalks at the same time, but by a process akin to amplifying the size of the needle itself.

This is an exciting time in quantum computing. Prototype small-scale quantum computers now exist in labs around the world, based on various physical implementations of the required quantum components. The most advanced of these are based on superconducting qubits [5,6] and trapped atoms [7] now reaching the 50 qubit level and beyond. With the ability to interface and program quantum computer hardware through cloud-based systems [8], the era of quantum software and application development is well and truly underway.

While the long-term vision of a universal (error corrected) quantum computer is reasonably well understood theoretically in terms of the types of tasks that could be carried out, the experimental and engineering challenges in realising such a machine pushes the expected realisation horizon out to probably decades. Quantum algorithms (e.g. quantum version of BLAS) running on error corrected quantum hardware could in principle assist in ML applications, although in lieu of the hardware existing to test these algorithms and the implementation caveats imposed means their relevancy is still an open question.

In the short to medium term the quantum computer space will be dominated by intermediate scale hardware comprising 100's of physical qubits with little or no quantum error correction. Broadly, the key question is what applications will benefit from such relatively noisy intermediate scale quantum computers (NISQC) [9]? This is an area of intense interest with applications/algorithms already being developed for chemistry (Variational Quantum Eigensolver) and optimisation (Quantum Approximate Optimization Algorithm). For the ML space, an important question is what adaptations of the quantum hardware, and associated quantum software/algorithms, will be of most benefit for ML applications?

Quantum algorithms:
A number of approaches in the quantum machine learning space rely on the existence of quantum algorithms which may speed up linear algebra and/or sampling tasks [10,11]. As an example, consider the HHL quantum algorithm [10] for matrix inversion, a common task in ML problems. For the inversion of a NxN matrix required to solve a linear system of equations Ax=b the HHL quantum algorithm scales logarithmically in N (modulo caveats below). This is the usual prima facie argument that quantum algorithms could be useful in ML with exponential speed-up. However, this statement of relative algorithmic complexity does not address the comparison in actual runtimes as the required quantum hardware does not yet exist, and various caveats still need to be understood. Associated with the application of quantum algorithms to data-rich problems are subtle input/output issues that need to be addressed in order to fully appreciate their potential. On the input side, the classical data in question (i.e. the matrix in the case above) must be loaded into the quantum register, and presented as a quantum superposition. The issue of loading data into a quantum computer is well known, and proposed approaches such as QRAM [12] attempt to address this. On the output side, the HHL algorithm, for example, does not produce the solution vector x directly, but rather a quantum register with the components of x in superposition. As pointed out by several authors [13], extracting the solution vector from the output register could involve some N measurements and erode the exponential speed-up of the algorithm, although it may be possible to extract more efficiently some global features of x.

Quantum hardware for ML sampling and neural networks:
Distinct from quantum algorithmic approaches to sampling [11], the use of specifically designed quantum hardware is being investigated to accelerate difficult sampling problems, e.g. those encountered in training restricted Boltzmann machines [14]. For example, the D-Wave system is a "quantum annealer" comprising some 2000 superconducting qubits (with relatively short coherence times) with transverse Ising interactions that are tunable. Here the idea is to represent the problem as an equivalent thermal distribution over a (complicated and highly connected) Hamiltonian encoded in the qubits and couplings of such a quantum annealing machine [15]. At a given iteration, the system is initialised in a well-known state and adiabatically evolved to the Hamiltonian in question and physically sampled. Limited experiments, both

physically and theoretically, have shown that while the approach is promising in some instances, issues such as limited hardware connectivity and "freeze-out" of the distribution mean the efficacy of this approach is still an open question [16].

At the other end of the special purpose quantum hardware are proposals for fully quantum neural networks, i.e. a quantum Boltzmann machine [17]. It is not fully known what the advantages are over the classical approach, however, the pertinent observation here is that while ML requires non-linear effects, quantum mechanics is inherently linear.

Classical ML for quantum technology:

In the reverse direction, there are fairly obvious opportunities for conventional ML to assist the development and deployment of quantum technology itself. For example, the design and optimisation of complex control sequences and/or analysis of quantum measurement data lends itself to a machine learning paradigm, and there are a number of examples of this application already [e.g. 18]. Ultimately, it has been suggested that a quantum RBM might prove useful in resource intensive exclusively quantum data analysis tasks, such as quantum state tomography [19].

Summary:

As a relatively nascent area, this brief report has only touched on some of the points of interaction between quantum technology and ML. This is an exciting area of research, and while there are indications that quantum information approaches could enhance ML, the actual speed-ups and applications remain open, and await actual quantum hardware to conclusively test [1,2] and/or discover new paradigms for using quantum technology in the ML space. The development of new approaches to ML using quantum information may in fact dictate the direction of the quantum hardware development (and vice-versa). In the Australian context, research and development in quantum hardware is very well supported through the Australian Research Council Centre of Excellence Scheme and the National Innovation and Science Agenda [20]. Research and support in quantum software, specifically associated with quantum/ML (and more broadly for that matter) is more diffuse – increasing opportunities for quantum software/hardware and ML communities to work together is an obvious pathway.

Acknowledgements:

The author acknowledges input from S. Tonetto, G. Mooney and C. Hill.

References:

[1] C. Ciliberto et al "Quantum machine learning: a classical perspective", Proc. R. Soc. A, **474** 20170551 (2018)
[2] J. Biamonte et al, "Quantum machine learning", Nature **549** 195 (2017)
[3] www.economist.com/news/essays/21717782-quantum-technology-beginning-come-its-own
[4] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe and J. L. O'Brien, "Quantum computers", Nature **464** 45 (2010)
[5] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets", Nature **549** 242 (2017)

[6] R. Barends et al, “Superconducting quantum circuits at the surface code threshold for fault tolerance”, Nature **508** 500 (2014)
[7] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z. X. Gong, and C. Monroe, “Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator,” Nature **551** 601 (2017)
[8] www.research.ibm.com/ibm-q/
[9] J. Preskill, “Quantum Computing in the NISQ era and beyond,” arXiv:1801.00862 (2018)
[10] A. W. Harrow, A. Hassidim, and S. Lloyd, “Quantum algorithm for linear systems of equations,” Phys. Rev. Lett. **103** 1 (2009)
[11] A. N. Chowdhury and R. D. Somma, “Quantum algorithms for Gibbs sampling and hitting-time estimation,” arXiv:1603.02940 (2016)
[12] V. Giovannetti, S. Lloyd, and L. MacCone, “Quantum random access memory,” Phys. Rev. Lett. **100** 1 (2008)
[13] S. Aaronson, “Read the fine print,” Nat. Phys. **11** 291 (2015)
[14] S. H. Adachi and M. P. Henderson, “Application of Quantum Annealing to Training of Deep Neural Networks,” arXiv:1510.00635 (2015)
[15] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, “Estimation of effective temperatures in quantum annealers for sampling applications: A case study with possible applications in deep learning,” Phys. Rev. A **94.** (2016)
[16] J. Marshall, E. G. Rieffel, and I. Hen. "Thermalization, Freeze-out, and Noise: Deciphering Experimental Quantum Annealers." Physical Review Applied **8** 064025 (2017)
[17] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko, “Quantum Boltzmann Machine,” Phys. Rev. X **8** 21050 (2016)
[18] S. S. Kalantre, J. P. Zwolak, S. Ragole, X. Wu, N. M. Zimmerman, M. D. Stewart, and J. M. Taylor, “Machine Learning techniques for state recognition and auto-tuning in quantum dots,” arXiv:1712.04914v2 (2017)
[19] Kieferová, Mária, and Nathan Wiebe. "Tomography and generative training with quantum Boltzmann machines." Physical Review **A96** 062327 (2017)
[20] https://www.industry.gov.au/strategies-for-the-future/boosting-innovation-and-science