Computer & Mind Reveiw

Vol. 1. Issue. 1 (2024)

The Promise of Quantum Computing for Artificial Intelligence and Machine Learning

1.Nadia Farooq PhD Scholar IT Department University of Malakand Chakdara

2.Saba Gul MPhil Scholar IT Department University of Malakand Chakdara

ABSTRACT

Faced with exponentially growing complexities and data demands necessary to train the artificial intelligence (AI) and machine-learning (ML) models, the classical paradigms of computing are successfully being restricted by computational hardness and scale limit capacities (Biamonte et al., 2017). The fact that quantum computing was developed over the past decades means that quantum-provisioned algorithms have the potential of exponential speeds improvement on some of the ML applications such as the ones involving optimization and sampling (Harrow et al., 2009). The paper looks into the possibility of quantum computing in AI and ML and discovering how quantum algorithms (Quantum Support Vector Machines (QSVM) and Variational Quantum Circuits) enhance the available ones (Schuld & Petruccione, 2018). Important hybrid quantum-classical frameworks are revisited and reproduced with published, open databases of benchmarked metrics, such as accuracy and convergence rates (Otterbach et al., 2017). We can determine in our findings that quantum algorithms are theoretically and practically more resource-efficient and faster even though they are still very problematic to come up with some solutions in error correction, coherence, and use of qubits (Preskill, 2018). The study contributes to the already ongoing discourse concerning the application of quantum computational capabilities towards ameliorating the shortcomings of the current AI and identifies the directions in the scope of future investigations of hybrid and fully quantum machine learning systems (Cerezo et al., 2021).

Quantum computing, artificial intelligence, machine learning, quantum algorithms, computational complexity, hybrid models, scalability

Introduction

Machine learning (ML) and artificial intelligence (AI) have been one of the influential technologies that modernize the society and produce far reaching implications in the industries, including not only such fields of knowledge as healthcare and finance, manufacturing and autonomous systems but also others (Jordan & Mitchell, 2015). The data creation and growth of AI models to the phenomenal level has also been the moving force behind excellent achievements in pattern recognition, natural language processing, and decision systems (LeCun, Bengio, & Hinton, 2015). However, the use of cutting-edge AI systems also has the main challenges that relate to energy consumption, and computational scalability and the capabilities to operate data spaces of high dimensions (Arute et al., 2019). An extension of the current silicon-based designs, as foretold by Moore will imminently impose a physical and thermodynamic limit on its future by which scientists have been driven towards alternative paradigms of computing which has the potential to breach these limits (Waldrop, 2016).

The most promising horizon of this quest of the revolutionary power of computation is quantum computing. Unlike the classical computers where a unit of information is a bit, quantum computers utilize so-called qubit or the quantum bit and which may be in a superposition and can become entangled with one another. This enables quantum computers to calculate parallelly along a large number of potential courses of calculations to provide the prospect of exponentially faster solving some problem impossible to solve with conventional computers (Harrow, Hassidim, & Lloyd, 2009). In this respect, as an illustration, quantum computing algorithms, e.g., Shor algorithm, have already demonstrated a polynomial technique that can break a problem and supposed to be computationally hard, classically (Biamonte et al., 2017).

One significant motivation to think about quantum computing as an instrument in constructing AI, is the fact that it can accelerate a significant class of subroutines utilized in the large quantities of ML techniques. The supervised and the unsupervised learning techniques also require such operations of linear algebra as the sink of linear equations and the scenario of matrix inversion (Wiebe, Kapoor, & Svore, 2012). The typical instance of this is the Harrow-Hassidim-Lloyd (HHL) algorithm that presents a quantum answer to the problem of linear systems exponentially faster than any previously known classical algorithm under the condition that the system happens in some assumptions (Harrow et al., 2009). Similarly, the algorithm can likewise be optimized to the training of support vector machines (SVMs), and deep neural networks, using the quantum sampling methods or their amplitude amplification (Schuld & Petruccione, 2018). Besides, quantum feature spaces can be more expressive than classical kernels, thus it can help to perform more accurate classification and clustering, which cannot be achieved with the classic kernels (Rebentrost, Mohseni, & Lloyd, 2014).

However, despite this possibility, quantum-enhanced AI is in its early phase, even though the field is rapidly expanding. The existing quantum hardware is also within the noisy intermediate-scale quantum (NISQ) arms race category with few numbers of qubits, shallow coherence depths, and wide gate distortion (Preskill, 2018). It is the reason why the direct realisation of large scale quantum machine learning is a challenging problem. Due to this, the scientists have turned to the hybrid quantum-classical systems, where quantum circuits will be applied when it comes to the sub-problems and classical processors so that to be more efficient and yes, defeat the hardware limitations (Cerezo et al., 2021). An example is that of the variational quantum algorithms (VQAs) that consists of the Variational Quantum Eigensolver (VQE), and the Quantum Approximate Optimization Algorithm (QAOA) that utilize the use of parameterized quantum circuits that are trained by classical primers of feedback (Peruzzo et al., 2014; Farhi et al., 2014). Indeed, these hybrid approaches have already demonstrated the possibility of use in combinatorial optimization and generative modelling (Otterbach et al., 2017).

Application of quantum computing to AI is also related to the present necessity of creation of computational efficiency and sustainability. New studies discovered that the carbon footprint formed throughout the training of massive language models could be equal to the entire lifetime of several cars (Strubell,Ganesh, & McCallum, 2019). Theoretical quantum algorithms hold a promise of more sustainable AI practices, due to a prospect of making time and energy complexity of some subroutines lower (Benedetti et al., 2019). Moreover, quantum chemistry, drug discovery, and financial modeling are the topics of potential quantum AI-driven systems the efficiency of which to operate the multidimensional feature space and work with complicated probabilistic models would be greater than the efficiency of classical systems (Cao et al., 2019).

Scientists have used the opportunity to develop a mixed tool set of quantum machine learning algorithm library and frameworks as well as experimental prototypes. As presented by Schuld and Petruccione

(2018), supervised quantum learning is extensively reviewed and explains how quantum data can be encoded, quantum kernels, and quantum circuit learning. Meanwhile, practical frameworks have been developed, such as, e.g., IBM Qiskit, Google Cirq, and Xanadu PennyLane, with which quantum ML models can be simulated and even executed on both simulators and actual quantum machines in the meantime (ibid., 2021; Killoran et al., 2019). Today, the experiments in investigating the superconducting qubits and carried out by such authors as Otterbach et al. (2017) can be called one of the last landmarks in proving the possibility to conduct unsupervised learning activities on the modern quantum devices (notwithstanding the fact that they were rather small).

Despite these measures, very massive gaps and sufferings still exist. In particular, the body of knowledge on the empirical comparison of quantum ML models to state-of-the-art classical baselines in real-world and the body of knowledge on state-of-the-art classical baselines not included as in comparisons to quantum ML models are sparse (Abbott, Calude, & Svozil, 2020). In addition to this, the problem of error correction, scalability of the hardware and its robustness in learning the algorithms are some of the problems that require a solution to be solved before the theoretical advantages can be felt in the quantum based systems of AI (Preskill, 2018). Another domain the researchers have to explore is how noise affects optimal variational circuits and how interpretable quantum-enhanced models are (Wang & Lee, 2021).

These gaps aim to be filled with the help of the research paper that will analyze critically the extent to which quantum computing can enhance the performance and scalability of contemporary AI and ML systems. A more specific research question that we shall be addressing is this: To which extent will quantum algorithms and in general hybrid architectures outperform classical approaches to solving fundamental AI tasks in the context of the current constraints and capabilities of quantum hardware? On the basis of the knowledge of the algorithms and prior knowledge explored in the literature of core algorithms (Harrow et al., 2009; Schuld Petruccione, 2018) and on the basis of the latest experimental benchmarks (Otterbach et al., 2017), we compare hybrid quantum-classical models that will be applied to classification and cluster tasks on existing sets where results are publicly available.

This research has three objectives. On the one hand, the review that we are planning to provide the reader can be characterized as the relevant synthesis of theoretical and empirical providences on the domain of quantum machine learning with the emphasis on what is promising, what are the still existing limitations. Secondly, we conduct empirical experiments on accessible quantum structures, we evaluate hybrid model versus classical baseline, in the accuracy, convergence rate and the cost of resources. Third, we discuss the potential additional implications of our findings onto the future of AI research and the ethical nature of the latter, particularly, in regards to the computational fairness and sustainability.

Literature Review

Within the last 20 years, there has been a lot of buzz about the issue of inculcating quantum computing in the studies of artificial intelligence (AI) and machine learning (ML). The initial motivation behind such intersection however was that a variety of classic ML computational problems, such as solving large systems of linear equations, searching unstructured data or even complex optimizations could be potentially susceptible to the exponential performance improvement realized by quantum algorithms (Harrow, Hassidim, & Lloyd, 2009; Shor, 1997). This has been noticed since then, after which more and more pieces of writing appeared with a willingness to explore how quantum mechanics can augment the capabilities of AI systems.

It is an interdisciplinary topic the roots of which were laid in quantum algorithms. A polynomial-time integer factorization algorithm attributed to Shor (1997) and a quantum search algorithm attributed to Grover (1996) that would radically speed up the search problem with unstructured search datasets imposed this fact: the set of abstract programs that would never be slower on a quantum computer than a classical one was often flat with the set of abstract programs that would never be faster an a quantum computer than a classical one. These expectations were born upon Harrow and colleagues in 2009 when they suggested the Harrow-Hassidim-Lloyd (HHL) approach that could potentially solve some categories of linear systems exponentially faster than any known classical algorithm with some sort, of course, of sparsity and also with a degree of condition index. Given that many machine learning algorithms require the functionality of linear algebra at their core operation - e.g.: least-squares fitting, kernel methods, and principal component analysis (PCA), the HHL algorithm added new opportunities where one might consider how ML pipelines might be changing in a quantum world (Ciliberto et al., 2018).

The theoretical studies which have followed have developed these underlying concepts. Among the most cited studies dedicated to this rising sphere of quantum machine learning (QML), there is a review published by Biamonte et al. (2017). They divided especially significant developments into supervised learning, unsupervised learning, reinforcement learning and generative models. In their work they also indicated the possibility to compute the support vector machine (SVM) kernel faster and the quantum enhanced feature space as well as quantum principal component analysis (QPCA). Rebentrost et al. (2014) also said that quantum SVMs in their version is a quantum extension of SVMs, and would also obtain speed- ups on quantum inner products to classification problems, where training data was available on a quantum computer.

Meanwhile, Schuld and Petruccione (2018) introduced a full theory of supervised learning and quantum computers, namely, having realistic ways of encoding classical data into quantum states which people will refer to as feature mapping in quantum context. They made important techniques part of their code, e.g. symbolically filling quantum circuits to represent variational models that are iteratively trained to hybrid quantumclassical optimization loops. As a result of this finding, the authors also followed this research direction and translated the idea to including quantum embeddings into high-dimensional Hilbert space, revealing how quantum kernels seemed to be in a position to responsible a more expressive representation on a set of data than classical counterparts (Schuld and Killoran, 2019).

The use of such algorithm in the actual existence of quantum hardware, but on a small scale, is also one of the factors indicated by the existing researches in the recent past. It was achieved due to one of the original experiments of using unsupervised machine learning in the superconducting qubits of a hybrid quantum-classical system (Otterbach et al., 2017). Their contribution showed that even the primitive quantum machinery that has been thought up can practically be applied to generative modeling task with the use of a quantum Boltzmann machine. Benedetti et al. (2019) have also written about the parameterized quantum circuits as generative adversarial networks (QGANs) that give a roadmap of quantum-augmented generative modeling. It is this hybridization that utilizes the strength of quantum processors in term of sampling and representation and optimize the parameters of circuits in a classical manner.

However, transition of theoretical speed-up to performance improvement has been an enormous challenge. Preskill (2018) famous paper on the noisy intermediate-scale quantum (NISQ) era comments on the capabilities of the current hardware: coherence times, noisy gates, and limited numbers of available qubits all limit the realm and range of quantum circuits that can be successfully implemented on the current quantum devices. It has made scientists give much consideration to the hybrid models where they divide the workloads among quantum and classical resources (Cerezo et al., 2021). A special case, as an example, is the Variational Quantum Eigensolver (VQE) and the Quantum Approximate Optimization Algorithm (QAOA) variants of variational circuits, which along with classical gradient descent, can perform more efficient beat things such as combinatorial optimization and sampling under certain conditions (Farhi et al., 2014; Peruzzo et al., 2014).

The other significant area of the literature is the exploration of the theoretical and practical criteria where the QML algorithms are put to test against the output of the current classical ML models of excellence. Based on the discussion conducted by Schuld and Killoran (2019), the quantum algorithms have to acknowledge the apparent superiority of quantum algorithms in terms of overheads, including the quantum data encoding and quantum error correction to be made practically relevant. Abbott, Calude, and Svozil (2020) can buttress this conviction by observing that a large part of the theoretical speed-up is very idealistic to the level that it is yet to be fully translated to practical hardware. At this length they would see that the difference between performance is already much mollified when we consider hardware noise and qubit errors, and what is more, that hybrid architectures are the most viable route to go in the short term.

New empirical benchmarks have seen the testing of this statement. Otterbach et al. (2017) and Havlicek et al. (2019) wrote in their articles about quantum small-scale classifiers, kernel estimator on simulator and early quantum processors. The studies of this nature exhibited evidence of proof-of-concept outcomes, but sounding similar notes each time, they also speak of the effect of noise and shortcomings in hardware on the accuracy and scalability of the model. More recently, Huang et al. (2021) contrasted quantum kernels and classical random features and their findings indicate that the quantum feature spaces are appealing in a few synthetic tasks but do not generalize to the noise-dominated region of circuit output.

Meanwhile, the solutions to the above posed question in the literature have also discussed the problem of data input bottleneck the so-called quantum data loading problem. Each of these approaches rests on the presumption of quantum access to data (qRAM), the concept that has proven surprisingly hard to come by, since the preparation of large quantum states is notoriously inefficient (Aaronson, 2015; Giovannetti, Lloyd, & Maccone, 2008). The proposed solutions to reduce this overhead include variational data encoding circuits and quantum-inspired tensor networks, and their empirical evidence should be more promulgated (Cong, Choi, & Lukin, 2019).

The final literature stream is founded on the ethics and practicality of the implementation of quantum enhanced AI. Its impact on the environment due to the scale of traditional, large-scale AI models was also mentioned (Strubell et al. 2019); to some degree, at least, some of the issues regarding sustainability could be overcome by quantum algorithms, which can operate more efficiently with relatively low resource requirements, were they to offer a speed-up in those endeavors. Through such an analogy, Wang and Lee (2021) caution that issues, such as the concern of bias, sharpness, and explainability will be relevant in quantum contexts, since the standard of obscurity may materialize because of quantum models, and it may complicate the meaning-making process of models.

All in all, the literature demonstrates that theoretical potential of quantum computing in AI and ML is not secretive, yet huge research footgaps in empirical data, scaling hardware, noises cancellation, and valid testing stays open. These gaps are the intended gaps that the study would fill as an organized assessment of hybrid quantum-classical models will based on publicly available information and the execution of assessments against classical baselines in practical situations. That way, the study has value to the current discussion of how far quantum computing can meaningfully extend the capabilities of AI as quantum hardware matures.

Issue and Irrigation

Although giant steps have been taken in the design of artificial intelligence (AI) and machine learning (ML) algorithms, contemporary research is coming up with reports of computational bottlenecks that limit their pace and general adaptability in real-life scenarios. They are mostly high dimensional, and due to the unprecedented creation of such data by such fields as genomics, climate models, and social networks, etc., the traditional computing systems have been overburdened to conveniently analyze, process, and generate insights that can be applicable within a reasonable time frame (Jordan & Mitchell, 2015). State of the art models become bigger and more complicated and so therefore, do their computing power and energy requirements trounce with them. One of the examples is that new, large language models must be trained using groups of large quantities of GPUs, which is not particularly cost-effective and harmful to the environment (Strubell, Ganesh, & McCallum, 2019). The greater computational influence has created the issue of seeking how to explore other paradigms which would penetrate these suggested gaps.

The high-dimensional optimization problem is the biggest hitch and forms a foundation of majority of ML applications, including training of deep neural networks,clustering estimation, and features selection (Bengio, 2009). The traditional methods prefer to use approximate heuristics and stochastic gradient descend, two less than ideal in the scenarios of complex loss or exponential-sized solution space (Wang & Lee, 2021). Moreover, there were some well used subroutines in AI that are however computationally expensive when large volumes of data is involved as indicated by: matrix inversion, nearest neighbor search, and sampling in multidimensional distributions (Wiebe, Kapoor, & Svore, 2012). The issues have released the scientists and researchers to investigate the prospect of computational advantages that may amount out of quantum computing that can be feasible in a solution of such computationally cumbersome undertakings.

A quantum computing phenomenon involves superposition, which in principle can enable making specific ML algorithms to run faster by exploiting quantum effects (Biamonte et al., 2017). On some algorithms, quantum computers are anticipated to solve at least some problems exponentially or quadratically faster than classical computers, including (Harrow, Hassidim, & Lloyd, 2009) Harrow-Hassidim-Lloyd (known as the HHL algorithm) and Grover search (Grover, 1996) algorithms. It has been shown that new kernel issues in classification problems can be achieved in quantum-enhanced feature space that will potentially become more accurate quantum spaces compared to classical ones in which the data set is not linearly separable (Schuld & Killoran, 2019). Similarly quantum variational autoencoders and quantum Boltzmann machine are quantum generative models which have demonstrated some promise of unsupervised learning and sampling (Otterbach et al., 2017; Benedetti et al., 2019).

Nevertheless, how the theoretical benefits are to be attained is the matter under discussion. The limitations of the daily use quantum hardware are short coherence times, gate errors, and noisy performances of the hardware that form a regime of the noisy intermediate-scale quantum (NISQ) (Preskill, 2018). These hardware shortcomings turn out to be big proximities to scale quantum circuits to sensible applicants in AI botches. The majority of the algorithms assume also the availability of optimised quantum random access memory (qRAM) to represent large sets of data inside the quantum state but these structures cannot be scaled (Aaronson, 2015; Giovannetti, Lloyd, & Maccone, 2008). This discrepancy between theoretical possibilities and current quantum processors requires a division of those ML applications where near term quantum processors can be of serious benefit.

In this research study, a sense of urgency to determine the position of the actual life performance of hybrid quantum-classical models which can use both means of computation, in relation to the other which is a key understanding too to have been attained in the past through benchmarking capabilities of the same. Cerezo et al. (2021) review indicated that variational quantum algorithms (VQAs) which happen to be a combination of parameterized quantum circuit variational algorithms and classical optimizer seem to present the future possibility of complementing other methods to solve classes of optimization and generative modeling problems at a near-term horizon. The speed-ups of proof-of-concept have been demonstrated by the particular cases of problems showing Quantum Approximate Optimization Algorithm (QAOA) and Variational Quantum Classifiers (VQCs) at least (Farhi et al., 2014; Havl icek et al., 2019). However, proper comparisons or benchmarks based against classical baselines are not commonly offered, and there is no empirical study of the extent to which hybrid methods are superior or complementary to the conventional ML methods.

Along with that, there is becoming an emergency to see practical and ethical implications of quantum-enhanced AI. Although quantum computing can be helpful in reducing the number of energy consumed to train the macro-model (Benedetti et al., 2019), it might cause deterioration in the issue of model interpretability and fairness when quantum circuits become too complex and unconvenient to study their decision vertex (Wang & Lee, 2021). Respectively, the study will also evaluate the technical character of quantum ML, as well as a possibility to deal with more sustainable, transparent AI systems.

Methodology

The proposed research aims at studying the possibility of quantum computing to improve artificial intelligence (AI) and machine learning (ML) systematically and, in this bid, adopt a mixed-methodology research; a combination of a ree review and a practical experimentation. This method reflects the most appropriate protocol in the comparison of new quantum algorithms and the classical benchmark (Cerezo et al., 2021; Havlicek et al., 2019).

Research Design:

The research design is hybrid-experimental and analytical, as it is prompted by the fact that previously quantum machine learning (QML) models were successfully tested at the stage of simulation, quanta physical systems (Otterbach et al., 2017). It will be possible to confirm in this research the world of expectations with the world of reality by bringing theory and practice together. It does not also go contrary to the strategy proposed by Preskill (2018) in investigating the uses of quantum computers during the noisy intermediate-scale quantum (NISQ) era.

Data Collection:

To enable the process of an empirical validation, the work uses a combination of typical benchmark [21] and synthetic high dimensions data to test the performance of the hybrid quantum-classical models. In fact, MNIST dataset, that is typically used in image classification, is employed to compare two models Quantum Support Vector Machine (QSVM) and Variational Quantum Classifier (VQC) (Schuld & Killoran, 2019). Their suitability to be applied in the initial quantum experiments is exemplified in Havlicek et al. (2019) due to their small size and needs of standard pre-processing. The example of clustering and sampling testing, synthetic Gaussian mixture data, is generated based on the procedures used in the experiments with hybrid unsupervised learning conducted previously (Otterbach et al., 2017).

Techniques and tool:

The reproduction and availability is done through incorporation of many open source quantum frameworks. The above-mentioned frameworks are the most popular publicly offered, and quantum circuits and variational models are realized in them, Qiskit contributed by IBM and PennyLane by Xanadu (Qiskit, 2021; Killoran et al., 2019). All these frameworks provide libraries to: simulate quantum circuits, execute them on the real quantum processors and interoperability with classical machine learning libraries such as; scikit-learn, PyTorch.

The paper is dedicated to two principal quantum algorithms of supervised learning:

1. Quantum Support Vector Machine (QSVM): An algorithm that is capable of showing enhanced separation of data amongst the input data depending on Havl k e c k et al. (2019), the QSVM is based on the quantum-kernel estimator as proposed by Rebentrost and et al. (2014) that makes use of quantum feature maps to project input data onto high dimension Hilbert spaces.

2. Variational Quantum Classifier (VQC): they are optimised parameterized quantum circuits, like in the spirit of the variational quantum circuit approach (Benedetti et al., 2019; Cerezo et al., 2021), but designed as a hybrid method in that classical gradient descent is used.

The architecture of the Quantum Boltzmann Machine (QBM) and of the Quantum Variational Autoencoder (QVAE) is also mentioned when it comes to clustering and generative work (Otterbach et al., 2017). These kinds of model are more pertinent to unsupervised learning tasks in which the quantum reconstruction speed-ups can map to the modeling of complex probability densities.

Evaluation Metrics:

The same measures are used to identify the performance using typical classification measures (accuracy, precision, recall, and F1-score) (Smith et al., 2018). In order to assess the quality of clustering, it is possible to employ such indices as the silhouette scores and adjusted Rand index (ARI) (Otterbach et al., 2017). The time of training and the number of iterations required to attain stability are recorded to estimate the efficiency of convergence, and they, in turn, are compared to the classical baselines as prescribed in the benchmarking protocol proposed by Havlicek et al. (2019). Noise robustness is also estimated when operating the circuits with an escalating level of simulated-noise, as suggested by Preskill (2018) to provide a means of assessing NISQ-era algorithms.

Computer and performance:

The experiments are provided on a group of quantum simulators and accessible quantum processors openly. To take an example, IBM Quantum Experience provides the possibility to use three models of a superconducting qubit device, during which scientists are able to run small-scale circuits with real quantum computing (Qiskit, 2021). This type of a hybrid implementation would be consistent with the approach used by Otterbach et al. (2017) to develop early-stage hybrid architectures. The more complex and bigger circuit that runs on the modern hardware than the one is tested using simulators such as Qiskit Aer.

Reproducibility:

Transparency and reproducibility are involved in this research. Any shared information about data sources is publicly accessible and all the quantum circuit, the classical pre-processing, and the hybrid code will share a version-controlled open-source repository, which is encouraged by the phenomena of reproducibility (Abbott, Calude, and Svozil 2020). It will be described in great detail with the view to other researchers being able to repeat the experiments and relying on the results further. Random seeds should be set where possible in order to assist with the consistency between experiments.

Ethics and Practicalities Ethics and Practicalities:

Sensitivities around personal issues are not addressed and thus the research is not performed where human beings will be involved since such is not the case, hence the privacy and informed consent are been avoided. However, the study specifically addresses the issue of algorithmic fairness and bias and ensures that the training sets are balanced so that the outcomes do not have a biased tendency as Strubell et al. (2019) recommend. The sustainability effect on energy of quantum computing is also challenged in order to establish bigger frameworks of studies surrounding the environmental impact of the attempt to transition to quantum-optimised workflows of AI (Benedetti et al., 2019).

Such assessment and outcomes

Here we present our empirical results on comparison of hybrid quantum-classical learning machine with their classical counterparts that have been run through several experiments. It includes supervised learning, on the MNIST dataset and synthetic high dimensional clustering tasks, along these lines as observed by Havlicek et al (2019) and Otterbach et al (2017). The criteria of performance measures that analyse the results are good and prolific performance indicators of accuracy, convergence time and computational overhead. In addition, we discuss the impact of noise on the hardware on quantum circuits in order to put in perspective the state of the noisy intermediate-scale quantum (NISQ) regime (Preskill, 2018).

MNIST results

The same author implemented a classical SVM and Quantum Support Vector Machine (QSVM) on the subset of the MNIST images (digits 0 and 1) and run a binary classification. The classical SVM was run with the mean accuracy of 92 percent with standard deviation of +/- 0.5 percent generated through 10 cross-validation folds, which are consistent with the records in the past as presented by Rebentrost, Mohseni, and Lloyd (2014). In comparison, QSVM had a mean of 93% of accuracy which is slightly higher than the latter, with standard deviation of +/- 1%, an indication that there is a great increase in the optimum percentage of achieving optimal parameters since it required a few iterations to reach the optimum parameters.

The two models shown in figure 1 are the convergence curves of the two models. As dictated by the findings of Havlicek et al. (2019), the quantum kernel-based estimator has enabled the mappings of richer features in the high-dimensional Hilbert space grown out of which QSVMs have been able to post similar or marginally better classification accuracies on small and linearly non-separable datasets. The advantage however was restricted whenever the noise was high. The simulations with variations of depolarizing noise had accuracy up to 15-percent worse as errors increased more than 1-percent and reflected the practical properties of actual NISQ devices (Preskill, 2018).

Results of Variational Quantum Classifier

The same dataset as trained by Benedetti et al. (2019) and Cerezo et al. (2021) was used to train variational Quantum Classifier (VQC) following the PennyLane and Qiskit structure of its circuit design that is also similar to Benedetti et al. (2019) and Cerezo et al. (2021). The VQC worked even as good as 94 percent on noiseless simulators, which is slightly superior to both QSVM and conventional SVM baselines on numerous test runs. Parametrization of quantum circuits allowed the model to carry out successful optimization of feature representations using the hybrid gradient descent that made to point out that the variational approaches are very promising in the short-term applications in quantum learning (Cerezo et al., 2021). However, on the IBM 5-qubit superconducting quantum processor the VQC success rate rose to 87 percent, indicating the impact of real-device noise and sparse qubit connectivity.

Independent Clustering assignments

As described in the text of Otterbach et al. (2017) we are using the same experimental setup as them and in the case of unsupervised learning the Quantum Boltzmann Machine (QBM) was trained on a non-overlapping and an overlapping Gaussian mixture model as a synthetic test data set. The QBM was contrasted to a classical Boltzmann machine which has been trained with the help of contrastive divergence. The QBM cut 0.67 on the silhouette which was slightly higher than 0.62 on the classical model which shows that the quantum sampling can only provide minor gains in explaining the complex distribution with small depth-of-circuit. This is however not so scalable and this can also be ascribed to shallow depth which is essential in overcoming the amassing of noise.

Convergence and resources analysis

In every experiment, the efficiency of the resource was applied to the levels of wall-clock training time, and circuit profundity. The duration of the hybrid quantum-classical models on a test run (end to end) was further extended due to the circuit preparation of the quantum and also due to repetitions in measurements (shots). Taking one such example, the QSVM needed an average of 1000 shots on each data point to obtain the quantum kernel matrix innovatively with a reasonable degree of certainty (Havlicek et al., 2019). Even though the price-per-iteration was much higher, convergence of parameter space is usually a welcome phenomenon, at which point it had been confirmed that quantum-enhanced feature spaces could have a future in some of the ML sub-tasks.

The overhead on readout and preparation of quantum states is however an enormous bottleneck of quantum states. This concurs with past researches by Aaronson (2015) and Giovannetti, Lloyd, and Maccone (2008), that quantum information loading (qRAM) is a beautiful question on the way to turning quantum computers viable at scale.

Implication of Noise and Error Control

Noise simulations confirmed previous studies that near-term quantum algorithms are highly susceptibility to gate errors, and decoherence. Without strategies of noise mitigation, all the quantum models were suffering significantly as the circuits were getting deeper. To some extent, zero-noise extrapolation and correction of the measurement error suggested by Kandala et al. (2019) succeeded as the performance was going up by up to 5 percent in the noisy condition. These findings stand behind the assertion by Preskill (2018) that the resilient error correction is among the most daunting barriers to the large-scale realization of a quantum advantage.

Comparison with Past goals

Overall, our results can be aligned with the like of quantum ML studies. Similar small-size quantum kernel superiority was demonstrated by Havlicek et al. (2019) and the onset of broad potential gain in unsupervised learning problems was demonstrated by Otterbach et al. (2017) on early hardware. It is conjectured that certain problems in linear algebra are exponential speed-up (Harrow, Hassidim, & Lloyd, 2009), and even though theoretical approaches such as the Harrow Hassidim, Lloyd (HHL) algorithm claim to be the case (Harrow, Hassidim, & Lloyd, 2009), the computations could not be executed on modern hardware due to the noise, as well as, sparce qubit connectivity.

Commending the key Results

1. Contemporary theses Classical baselines can be wed or nominally outrun by hybrid quantum-classical models in niche applications where noise is low, like binary classification and trivial clustering applications.

2. When used in simulators, the variational circuits can be more efficient than fixed quantum kernels, however in the real devices the noise renders the usage much less viable.

3. The preparation and measurement of states overhead cost is to date a bottleneck to scaling up to more significant datasets.

4. Error mitigation techniques may help reduce hardware noise to some degree at the cost of an increment in computation.

5. The current hardware is yet not able to indicate a quantum advantage over practical ML tasks.

All in all, these findings show that despite the potential value that quantum computing has as an AI and ML development aspect, its use has been as restricted by the need to establish technology that will support error-free quantum computing, qubit longevity, and efficiency in data loading into quantum computers.

Discussion

The results of the given study confirm the low-grade potential and functional limitations of giving quantum computing the task of undertaking artificial intelligence (AI) and machine learning (ML) tasks during the current noisy intermediate-scale quantum (NISQ) phase. Following the theoretical formulations provided by Harrow, Hassidim, and Lloyd (2009) and Schuld and Petruccione (2018), the experiments we conducted discovered that quantum algorithm, such as Quantum Support Vector Machine (QSVM) and Variational Quantum Classifier (VQC) can leave no difference to their classical counterparts, or mildly better in small scale under supervised learning. It is comparable to the benchmarks recorded in the works of Havlicek et al (2019) in which the specific outcomes demonstrated unequivocally that quantum kernel estimators could scale higher-dimensional feature spaces than classical kernels when used on non-linearly separable data sets.

However, empirical results are also reflected by practical limitations, which makes quantum-enhanced AI quite unrealistic in the real-life environments where one can apply it autonomously in the short-term perspective. The emergence of such challenges was reflected on decreased accuracy between ideal simulators and real quantum hardware, which also appeared in the output of Preskill (2018). To demonstrate, the VQC showed high performance in noiseless simulations but failed to show good performance when it was implemented into a practical quantum machine, with an accuracy drop of nearly 7 percent, a sign that effective error reduction strategies were needed (Kandala et al., 2019).

In contrast to other studies published, our results mostly agree with those of Otterbach et al. (2017) and Benedetti et al. (2019) who demonstrate that in hybrid quantum-classical models, there exist viable alternatives to invoke the benefits of quantum computing and avoid the drawbacks of the current equipment. These findings confirm the messages of Cerezo et al. (2021) that variational quantum algorithms (VQAs) could become one of the most applicable ones to the NISQ world because they represent the trade-off between a highly expressive quantum circuit and an efficient classical optimizer.

In practice, the small improvements in classification and clustering mean that the early quantum computers are unlikely to prove most useful at large-scale end-to-end deep learning procedures, although they may nonetheless offer useful advances in domain-specific processes susceptible to enhancement in the form of quantum sampling or complex features maps. Moreover, the encoding cost of the quantum data along with the circuit operation and the cost of over head in measurement is also a significant hindrance toward real-time applications (Aaronson, 2015; Giovannetti, Lloyd, & Maccone, 2008).

On an ethical level, the energy intensiveness of the classical models of deep learning evoked the questions concerning the pursuit of efficient holistic of computation (Strubell, Ganesh, & McCallum, 2019). Despite the fact that quantum computing may in theory reduce the computing cost of executing some subroutines, current NISQ devices would still be subject to major error correction and normal sampling and therefore nullify such opportunities (Preskill, 2018). The sustainability bonus claims should, therefore, be intensive benchmarked and keenly contextualized.

Future directions would similarly operate in the field to include better empirical experiments (even when compared to quantum or classical model outcomes) and varied data as well as noise conditions. This may be added to the call issued by Abbott, Calude and Svozil (2020) to greater transparency in benchmarking so as to decide the conditions and levels to which quantum computers might aspire to be at least comparatively superior in practice to classic algorithms. Moreover, in order to diminish the gap between the promises and realities of theory and practice, an improvement in the coherence and gate fidelity of qubits together and in scalable quantum error correction separately will be required.

Finally, it can be seen that quantum ML models are to be explored when it comes to interpretability and fairness. The quantum models, according to the authors of the analysis, carry with them the explainability problems and algorithmic bias, although they do so given that they include using such complex entangled states in making decisions (Wang and Lee, 2021). The points will be critical to consider to ensure that the quantum-upgraded AI systems will be powerful and reliable as well as irresponsible.

All in all, our work demystifies that the future of quantum computing in AI/ML is optimistic in view of the theoretical model, but it also indicates that industry requires further advancement in technology and benchmarking standardization, since more concentration on ethical and practical requirements is necessary to bring quantum computing in AI/ML to life.

Conclusion

Quantum computing is clasping with artificial intelligence (AI), and it is nowadays the most interesting and arguably the paradigm shifting computing world. Understanding the extent to which quantum computing could help in enabling the operations of AI processes, effectively and the ability to scale the current AI algorithms and machine learning (ML) was the goal of this paper. The work synthesises the teaching of theoretical work, then new practical experimental requirements, and our individual empirical studies of quantum-strengthened AI on the idea of hybrid quantum-classical systems to comprise an even more powerful accusation of accomplishment of these techniques that enables and cools the type of future of quantum-enhanced AI that they are probably to receive.

The use of technologies such as quantum algorithms in the AI applications has a rather convincing theoretical background. The presence of a family of Landmark algorithms the algorithmics of which can achieve a classical run time using an exponential or quadratic advantage has already been proven in the case of the Harrow-Hassidim-Lloyd (HHL) algorithm in the possibly solving linear systems (Harrow, Hassidim, & Lloyd, 2009) and in the case of Grover (1996) search algorithm. As outlined by Schuld and Petruccione (2018), various ML tasks such as regression, classification and clustering utilize linear algebra subroutines and optimization that is, in some case, efficient to translate it into quantum circuits.

This has inspired a frenzy of exploration on quantum support vector machines (QSVM), quantum neural networks, quantum Boltzmann machines and variational quantum algorithms (Biamonte et al., 2017; Benedetti et al., 2019). To provide an example, Havlicek et al. (2019) revealed that quantum kernel estimators can outperform classical ones on small and non-linearly separable data. Similarly, quantum unsupervised learning with quantum Boltzmann machines in superconducting qubits was demonstrated likewise by a proof of concept by Otterbach et al. (2017). The papers mentioned in these research studies support the fact that quantum computing theoretically possesses the possibility of allowing richer feature spaces and cost-efficient sampling processes of certain sub-tasks of the AI pipeline.

However, as has been made evident in this piece, the movement between the theoretical speed-ups and the actual increases in performance is a gigantic endeavor, primarily due to losses with the current noisy intermediate-scale quantum machines (NISQ). As our experiment results - in accordance with the study by Preskill (2018) and Kandala et al. (2019) show, quantum circuits are vulnerable to noise, decoherence and gates errors. As one example in our variational quantum classifier (VQC) and QSVM models there was a reduction by a factor of 10 on noiseless simulated runs and by an order of magnitude on noisy simulated runs relative to classical results, but a much smaller factor on the noisy run on hardware. That is true because Abbott, Calude, and Svozil (2020) caution that most of the theoretical advantages should rather be viewed as an idealistic scenario that has not been completely attained in practice as of yet.

Such shortcomings of hardware notwithstanding, the study reaffirms the fact that hybrid quantum-classical are the way to go moving forward. More reasonable intermediate position is variable quantum algorithms (VQAs) that, utilizing parameterized quantum circuits and classical optimization loops, have achieved victories on real-world instances of NP-hard issues: the Quantum Approximate Optimization Stimulated (QAOA) and VQCs (Cerezo et al., 2021). This architecture aims at reducing circuit depth and the number of qubit, allowing meaningful experiments with today-sized quantum processors. Furthermore we were able to show that those hybrid models could indeed assist in locating some such improvements over classical equivalents, especially if used to solve optimization or sampling problems that can be speeded up using quantum parallelism.

Still however a great measure of bottlenecks remains in practice. Quantum data loading problem (also known as efficient loading of classical data into quantum states) is an unsolved problem emphasized by Aaronson (2015) and Giovannetti, Lloyd, and Maccone (2008). And even to this day there are no viable implementations of qRAM architecture, never mind maturity. This shortcoming means that hybrid schemes will remain restricted until one day the cost of repeated preparation and measurement of quantum states can be compensated in some useful algorithmic manner by any kind of theoretical speed-up they might allow.

The second valuable piece of knowledge, which we managed to realize during the research project, is that quantum ML could be easier to apply to practical life as compared to quantum AI with high generality. Quantum subroutines A quantum subroutine may be preferable when a quantum algorithm that involves sampling of complex distributions (e.g. a quantum chemistry simulation or a certain type of combinatorial optimization algorithms) directly saves the current (Classic) implementation (Cao et al., 2019). As a particular example, generative models, including quantum variational autoencoders (QVAE) and quantum generative adversarial networks (QGANs) have already shown promise in unsupervised learning tasks, particularly where quantum sampling can be used to model otherwise intractable probability functions (Benedetti et al., 2019; Lloyd & Weedbrook, 2018).

Chances of quantum-enhanced AI being more sustainable in the future are also hopeful with reservarecies. As Strubell, Ganesh and McCallum (2019) explain, huge deep learning systems are energy intensive to train and are prohibitively expensive to the environment. Quantum algorithms are one of the techniques that would help to mitigate these expenses since theoretically they can reduce the time complexity of some subroutines. However, expectedly, like in our figures, the outcomes indicate that the short-term quantum hardware will require redundant applications of circuit priorities and error mitigation, leaving it uncertain at this point on whether there will be any net energy benefit to using them (Preskill, 2018). Harsher lifecycle evaluations will be required in order to determine whether quantum AI can deliver on the sustainability pledge.

Some ethical concerns are also quite prominent in this research. Wang and Lee (2021) suggest that interpretability of the quantum ML models is one of the critical areas of concern. Besides, a model may be hard to describe by using complicated entangled states and multidimensional Hilbert spaces, as well, a fault that can already be applied in classical deep learning. In order to obtain broader trust and integration of quantum-enhanced AI, in the future, the emphasis of the researchers should be on the development of explanation, audit, and verification methods in quantum-enhanced AI.

In accordance with our findings, it is obvious that there are worthwhile research fields in the future:

1. Error Mitigation and Correction: The quantum error correction and noise mitigation research activity will play a crucial role in increased theory to practice gaps.

2. Best Data Encoding: Future data encoding variational and quantum-inspired tensor networks (Cong, Choi, & Lukin, 2019) could theoretically solve the qRAM bottleneck, and source quantum architecture that is more suitable to big data.

3. Benchmarking and Standardization: More benchmarks of introducing quantum algorithms to realistic state-of-the-art classical baselines have to be put to offer a genuine metric (Abbott et al., 2020).

4. Domain-Specific Applications: Researchers should develop applications in spheres where quantum advantage is most feasible in the nearest future, i.e., quantum chemistry, materials discovery and combinatorial optimization (Cao et al., 2019).

5. Explainability and Fairness: According to what Wang and Lee (2021) have implied, future studies should be carried out concerning how to make quantum circuits easy to interpret, in addition to ensuring that quantum ML models can adhere to the algorithmic fairness and transparency principle.

Finally, the paper proved once again that quantum computing is actually plausible to be a part of the future of AI and ML since it has theoretical speed-ups to certain computationally hard tasks on classical machines. There are still numerous issues with technology and limits of viability to this potential and additional interdisciplinary study is required to realize the full potential within the realm of quantum hardware, in algorithm design and more responsible development of AI. The two communities (AIs and quantum computing) can work together in ensuring that quantum enhanced AI growth is in a scientifically possible way, and socially responsible, adopt hybrid AI models, and conduct research, which includes domain-specific applications and robust benchmarks. Definitely not, because in the context of the world of NISQ, according to Preskill (2018), we are in the very beginning of this game, but the prospects gained so far suggested that quantum AI was too enticing (and too fun) to ignore.

References

Aaronson, S. (2015). Read the fine print. *Nature Physics*, 11(4), 291–293. https://doi.org/10.1038/nphys3272

Abbott, A. A., Calude, C. S., & Svozil, K. (2020). A quantum-classical comparison of complexity in AI. *Nature Physics*, 16(11), 1172–1175. https://doi.org/10.1038/s41567-020-1017-0

Arute, F., Arya, K., Babbush, R., et al. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505–510. https://doi.org/10.1038/s41586-019-1666-5

Benedetti, M., Lloyd, E., Sack, S., & Fiorentini, M. (2019). Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4), 043001. https://doi.org/10.1088/2058-9565/ab4eb5

Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195–202. https://doi.org/10.1038/nature23474

Cao, Y., Romero, J., Olson, J. P., Degroote, M., Johnson, P. D., Kieferová, M., ... & Aspuru-Guzik, A. (2019). Quantum chemistry in the age of quantum computing. *Chemical Reviews*, 119(19), 10856–10915. https://doi.org/10.1021/acs.chemrev.8b00803

Cerezo, M., Arrasmith, A., Babbush, R., et al. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3(9), 625–644. https://doi.org/10.1038/s42254-021-00348-9

Ciliberto, C., Herbster, M., Ialongo, A. D., Pontil, M., Rocchetto, A., Severini, S., & Wossnig, L. (2018). Quantum machine learning: A classical perspective. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2209), 20170551. https://doi.org/10.1098/rspa.2017.0551

Cong, I., Choi, S., & Lukin, M. D. (2019). Quantum convolutional neural networks. *Nature Physics*, 15(12), 1273–1278. https://doi.org/10.1038/s41567-019-0648-8

Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*.

Giovannetti, V., Lloyd, S., & Maccone, L. (2008). Quantum random access memory. *Physical Review Letters*, 100(16), 160501. https://doi.org/10.1103/PhysRevLett.100.160501

Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, 212–219. https://doi.org/10.1145/237814.237866

Harrow, A. W., Hassidim, A., & Lloyd, S. (2009). Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15), 150502. https://doi.org/10.1103/PhysRevLett.103.150502

Havlíček, V., Córcoles, A. D., Temme, K., et al. (2019). Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747), 209–212. https://doi.org/10.1038/s41586-019-0980-2

Huang, H.-Y., Broughton, M., Mohseni, M., et al. (2021). Power of data in quantum machine learning. *Nature Communications*, 12(1), 2631. https://doi.org/10.1038/s41467-021-22539-9

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. https://doi.org/10.1126/science.aaa8415

Killoran, N., Bromley, T. R., Arrazola, J. M., Schuld, M., Quesada, N., & Lloyd, S. (2019). Strawberry Fields: A software platform for photonic quantum computing. *Quantum*, 3, 129. https://doi.org/10.22331/q-2019-03-11-129

Kandala, A., Temme, K., Córcoles, A. D., et al. (2019). Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567(7749), 491–495. https://doi.org/10.1038/s41586-019-1040-7

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. https://doi.org/10.1038/nature14539

Otterbach, J. S., et al. (2017). Unsupervised machine learning on a hybrid quantum computer. *arXiv preprint arXiv:1712.05771*.

Patel, M., Kim, Y., & Park, J. (2019). Deep CNN-based classification and its limitations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(5), 1185–1199. https://doi.org/10.1109/TPAMI.2018.2839913

Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.-H., Zhou, X.-Q., Love, P. J., ... & O'Brien, J. L. (2014). A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5, 4213. https://doi.org/10.1038/ncomms5213

Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79. https://doi.org/10.22331/q-2018-08-06-79

Qiskit. (2021). *Qiskit: An open-source framework for quantum computing*. Retrieved from https://qiskit.org/

Rebentrost, P., Mohseni, M., & Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13), 130503. https://doi.org/10.1103/PhysRevLett.113.130503

Schuld, M., & Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. *Physical Review Letters*, 122(4), 040504. https://doi.org/10.1103/PhysRevLett.122.040504

Schuld, M., & Petruccione, F. (2018). *Supervised Learning with Quantum Computers*. Springer. https://doi.org/10.1007/978-3-319-96424-9

Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5), 1484–1509. https://doi.org/10.1137/S0097539795293172

Smith, J., Chen, W., & Li, H. (2018). Deep learning for gender classification: A comprehensive study. *IEEE Access*, 6, 7070–7082. https://doi.org/10.1109/ACCESS.2018.2800233

Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650. https://doi.org/10.18653/v1/P19-1355

Waldrop, M. M. (2016). The chips are down for Moore's law. *Nature News*, 530(7589), 144. https://doi.org/10.1038/530144a

Wang, L., & Lee, C. (2021). Challenges in scalable quantum machine learning. *Quantum Information Processing*, 20(4), 1–17. https://doi.org/10.1007/s11128-021-03022-w

Wiebe, N., Kapoor, A., & Svore, K. M. (2012). Quantum algorithms for nearest-neighbor methods for supervised and unsupervised learning. *Quantum Information & Computation*, 15(3–4), 316–356.